Index of Swarm documentation (1.0.5)

documenting Swarm 1.0.5

This release of the documentation coincides with version 1.0.5 of the Swarm source code distribution. Though a complete manual for Swarm is still not available, this list of resources should give you a good place to start. Please direct all questions to the official Swarm mailing list: <swarm-support@santafe.edu> (mailto:majordomo@santafe.edu)
.

Table Of Contents:

Introduction

Our home page provides a good introduction to Swarm (what it does and what it aims to be). On the other hand, the apps-intro document aims to give you a more detailed overview of the sorts of things a user needs to do in order to get a simulation up and running in Swarm. The combination of these two documents should help you decide whether Swarm would be a useful tool in the context of your research.

http://www.santafe.edu/projects/swarm/
(http://www.santafe.edu/projects/swarm/)
: Swarm's Internet presence.

intro.html
(intro.html)
: What a Swarm application looks like, and how it is developed.

Installing Swarm

Swarm uses quite a few libraries and is intended to run under all major flavors of Unix. This presents the unpleasant but largely unavoidable side-effect of making Swarm hard to install. Ideally, you should get the sys-admin for your site to do the job. In any event, the following resources should prove useful:

http://www.santafe.edu/projects/swarm/release.html
(http://www.santafe.edu/projects/swarm/release.html)
: Swarm release web pages.

install.html
(install.html)
: Installation instructions

overview.html
(overview.html)
: Overview of the Swarm system.

To Do list
(todo.html)
: Additions we want to have to Swarm.

FAQ.html
(faq/FAQ.html)
: FAQ for the Swarm project.

LICENSE-libs
(LICENSE-libs)
: License terms for the distribution of the Swarm libraries.

LICENSE-apps
(LICENSE-apps)
: License terms for this distribution of the Swarm example applications.

Compatibility

1. 1.0.4 => 1.0.5

   - EZGraph's setGraphWindowGeometryRecordName and GUISwarm's setControlPanelGeometryRecordName have been retired. The macro SET_WINDOW_GEOMETRY_RECORD_NAME can now be used in any geometry archiving context.
   - After adjusting an application per previous item, be aware that the new internal naming conventions used by 'GUIComposite' classes (e.g. EZBin and GUISwarm) will probably differ from the archiving keys in your application. Such widgets will probably lose their saved placements.

2. 1.0.3 => 1.0.4

   - doTkEvents has been moved from ControlPanel to ActionCache. Be sure to direct all doTkEvent messages to the Swarm's instance of ActionCache and not the ControlPanel (a warning message will be generated otherwise).
   - The swarmobject library has been renamed to objectbase. Although the build procedure creates a link to maintain compatibility, applications should include <objectbase.h> instead of <swarmobject.h>.
   - Avoid calling globalTkInterp in all future applications, it will not be supported from version 1.1 onwards.
   - If there exists a <LIBRARY.h> file, e.g. <simtools.h>, that file is the advertised interface to LIBRARY, and it will defined in terms of protocols. Whenever possible, define variables in terms of the protocol they respond to:

   id <ProbeMap> probeMap;
   rather than using a static type:
   ProbeMap *probeMap;
   (One advantage of this is that fewer imports will be needed in your application.)

3. 1.0.2 => 1.0.3

   - Activity Library: Reference Release-Note #3: The AutoDrop option for any concurrent groups is automatically set to the same as the schedule in which it is contained. (Any existing setting of the option on the concurrent group type is ignored.) So, in 1.0.2, if you didn't set the AutoDrop flag for a concurrentGroup initialized under a schedule, it defaulted to NO or false.
   - Random Library: 13, 14

     a. Reference Release-Note #13 This release includes Random v0.7, written by Sven Thommesen. This version adds several new bit generators and distributions and rearranges the library. (Ref. $SWARMHOME/src/random/docs/WHATS.CHANGED.in.0.7 and WHATS.NEW.in.0.7)

     b. Reference Release-Note #14: With the addition of the ability to make use of the default random number generators and distributions such that different runs start with a different seed or with the same seed, the static seed was changed. This means that runs with 1.0.3 will not match runs made with the static seed for 1.0.2. This is acceptable because the use of the default generators and distributions is DEPRECATED.

4. 1.0.0 => 1.0.1

There should not be any imcompatibilities between 1.0.1 and 1.0.0. There are a couple of changes that affect the behavior of Swarm, however. The big ones are:

- The automatic dropping of probe displays upon the dropping of an object to which those probe displays were attached. This could break your application code if you leave in the dropProbeDisplaysFor: message where you drop such an object.
- The default probeMap has changed. This is only an issue when createProbeDisplayFor: is called without having previously created and installed a probeMap for that class. The new behavior is to create a probeMap on-the-fly that contains only the instance variables for that class.

See the Release Page

(http://www.santafe.edu/projects/swarm/release.html#changes)
for other differences.

5. *Beta* => 1.0.0

As can be expected with any software package, it is sometimes unavoidable that changes in the functionality of the package will cause incompatibilities with earlier versions. This is especially true when a package is a "proof-of-principle" package like Swarm. We've made an attempt to compile all the problems a user might have moving to the new release and put them here. Please read this thoroughly to decide what you might need to do to your app to get it to work with the new release.

a. The biggest and most pervasive problems will be due to the new *random* library. Notes on how to deal with this problem are provided in the documentation for the *random* library
(src/random.html)
.

b. A rather benign problem results from the repackaging of the *swarmobject* library. The interface to this library was brought into sync with the *defobj*and *collections* libraries. (The rest of the errant libraries will follow in a later release.) The solutions to the problems associated with this interface change are detailed in the documentation for the *objectbase* library
(src/objectbase.html)
.

c. There are a couple rather benign incompatibility introduced with 1.0 in the new *activity* library. Many parts of *activity* have changed. But, for the most part, everything works exactly the same. For details on the incompatibilities please see the documentation for the *activity* library
(src/activity.html)
. Briefly, the incompatibilities are:

- The high-level structural changes in the activity library has led to the renaming of the variable swarmActivity. It is recommended that [self getActivity] be used in its place (if your app even used this variable, it was probably in the *BatchSwarm*). This is a new message and it takes the place of the getSwarmActivity message. But, the obsolete message has been left in place for backwards compatibility.
- The getCurrentActivity() macro is gone. If you used the old getCurrentActivity() in your code, it won't work now. Use of this macro was not very widespread, since its main use is to access activity library internals. One of the other macros should be sufficient for any application.

d. The functionality of *Zone*s has been greatly improved and expanded upon. Most of the aspects of the idea behind *Zone*s are now in place. However, there is one incompatibility that must be noted in case your code is fairly old. The dropFrom: message has been removed. Even though it was still present in recent releases, its behavior was identical to drop. Any existing usage should be replaced by a simple drop message without any zone argument. SwarmObject subclasses are now restricted from accessing the zone that was once contained in an instance variable; the message getZone must be used instead.

Objective C Resources

The main technical skill required of a Swarm user is the ability to write Objective C code. Other recommended skills include the ability to use tools such as gdb and emacs. For the record, if you know C and have some form of experience in either C++ or Smalltalk, then learning Objective C should take no more than a day.

Objective C references
(http://www.santafe.edu/projects/swarm/ObjC/objective-c.html)
: Nelson's list of Objective C references on the Web. Includes a 10 minute overview to Objective C that is most of what you need to know about Objective C.

Objective C - Swarm Style
(swarmObjC.html)
: The Swarm system provides a few extensions to the syntax/style of Objective-C. This note describes those features which are important to know as a beginner and those which can be ignored...

Object-Oriented Programming and the Objective-C Language
(http://devworld.apple.com/dev/SWTechPubs/Documents/OPENSTEP/ObjectiveC/objctoc.htm)
: An excellent (online!) book on Objective C. This is for generic NeXT Objective C: Swarm uses GNU Objective C in addition to our own defobj (src/defobj.html) extensions.

debugging tips
(debugging.html)

When adding printf() just isn't enough...

emacs tips
(emacs-objc.html)
Some hints for editing Objective C code under emacs.

Core Libraries

Swarm is essentially a collection of libraries against which you link your simulation code. The following is a list of the those libraries which are crucial to the proper execution of Swarm, and are therefore shipped with the Swarm distribution.

- Defobj Library
(src/defobj.html)
- Collections Library
(src/collections.html)
- Activity Library
(src/activity.html)
- Objectbase Library
(src/objectbase.html)
- Space Library
(src/space.html)
- Random Library
(src/random.html)
- Simtools Library
(src/simtools.html)
- Graphical Interface Library (tkobjc)
(src/tkobjc.html)
- Analysis Tools
(src/analysis.html)
- Probe Interface
(src/probing.html)

Sample Application

The following sample applications represent the current best source of guidelines/advice about the structure of Swarm applications and the way in which they should be written (along with `intro.html`

(intro.html)
). In particular, the Heatbugs code has been thoroughly commented and demonstrates some of the main features in Swarm. The template application is a good place from which programmers can copy code as a starting point for their

initial simulations. There is more complete information (the screen shots are a bit dated) about a few of these demo applications in `../examples/` (`http://www.santafe.edu/projects/swarm/examples/index.html`)
.

- `Heatbugs, a basic complex system`
`(apps/heatbugs.txt)`
`: thoroughly commented!!!`
- `Template application`
`(apps/template.txt)`
- `Simple model of a stock market`
`(apps/market.txt)`
- `Mousetrap model of fission, discrete event simulation`
`(apps/mousetrap.txt)`
- `Tutorial is an app that progressively moves from a simple C implementation of a CA (with bug agents on it) up to a Swarm-style application.`
`(apps/tutorial.txt)`
- `Hello, World! is another simpler tutorial`
`(apps/hello-world.txt)`

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

`Last modified: 1998-01-08`

A Top-Down Introduction To Implementing an Experiment Using Swarm

---

This document attempts to explain the logical structure of a Swarm experiment application. Starting with a very general outline of an idealized experimental procedure, we successively increase the level of specification of each stage of this idealized structure until we arrive at details of an actual running Swarm application.

Along the way, we introduce and describe (very briefly) some of the tools currently available in Swarm to help users build experiments. The tools presented here are only suggestive: the Swarm library documentation and example applications will give a more detailed view of using specific components of Swarm.

These tools are more fully documented in the appropriate library directories. This high-level overview skips much of the detail in an effort to illustrate the logical structure of a Swarm application and to provide some motivation for why things are done the way they are in the example applications.

How we'll proceed

In the following, we will start out with a very high-level outline/abstraction of the "generic form" of an experimental procedure. Then we will increase the magnification on each piece of the outline until we reach specific details of Swarm code. If you compare the gross structure of the example applications you will see that despite the differences in detail, the overall architecture of Swarm applications adheres to the general abstraction of an experimental procedure.

The application we will use as an example is the Heatbugs application. The source code for that example is a thoroughly commented application and intended to be read by new Swarm programmers. It can be a useful template to start from: so can the very-stripped-down "template" application. We recommend that you gradually ease yourself into the Swarm perspective by experimenting with the example applications and branching out by modifying them in various ways until you catch on to the conceptual style of a Swarm Experiment.

Mag 1x: Experimental Procedure in a Computer

At the highest level of abstraction (= the lowest level of magnification), most experiments look like this:

1. Set up the physical system to be studied
2. Set up and calibrate the instrumentation
3. Run the experimental system and record the outputs of the instrumentation
4. Analyze results
5. If not happy with the results:
   - Change experimental and instrumental setup
   - Go to 3
6. Publish paper -> tenure -> fame -> etc.....

The important part of step 6) is that the published paper includes enough detail about the experimental setup and how it was run so that other labs with access to the same equipment can recreate the experiment and test the repeatability of the results. This is hardly *ever* done (or even possible) in the context of experiments run in computers, and the crucial process of independent verification via replication of results is almost unheard of in computer simulation. One goal of Swarm is to bring simulation writing up to a higher level of expression, writing applications with reference to a standard set of simulation tools.

First, let's look at what happens when we port the above stages into the world of a computer. In a computer, you don't just drag the pieces of your experiment in from the outside world and hook them up. You have to create a world with space and time, a bunch of objects in that world (stuff to study and stuff to look at it with), schedules of events over those objects, all sorts of computer widgetry to interact with that artificial world and to manage multiple experimental runs and the data that they generate, and so forth. In other words, in a computer, one usually has to first *create* from scratch all of the bits and pieces of the experimental setup - the virtual equivalent of beakers, bunsen burners, microscopes etc.

Perhaps the most important difference between an experiment in the "real" world and an experiment inside of a computer is the nature of time. In the real world, everything in one's experimental setup is moved forward in time via a very concurrency courtesy of the laws of physics. In a computer experiment, however, the experimenter has to explicitly move every object in his/her artificial universe forward in time, making sure that everything remains within some well-understood state of synchronization. Many fundamental problems in computer science have arisen in the course of trying to understand how to control and use concurrency. Furthermore, most people who implement computer simulations aren't even aware of the subtle, but quite-possibly dominating, impacts of assumptions that they aren't even aware that they are making about concurrency in their model when they code it up and run it.

Therefore, a very important aspect of setting up an experiment in a computer is how one weaves the multiple threads of time that must be woven together coherently in order to produce reliable, repeatable results. Much of our work on Swarm has been devoted to not only making the task of managing concurrency manageable, but towards mechanisms to make people aware that they are *always* making implicit assumptions about how multiple threads of time are interacting with one another in their experimental setups. Swarm forces experimenters to make their concurrency assumptions explicit, so that others can reproduce their results by implementing the same assumptions about the flow of time.

Mag 2x: Basis of Swarm Computation

Here is a first approximation to embedding the above outline of an experimental procedure in Swarm code:

Swarm is implemented in the Object-Oriented Programming language Objective-C. Computation in a Swarm application takes place by having objects send messages to each other. The basic message syntax is

```
[targetObject message Arg1: var1 Arg2: var2]
```

Where "targetObject" is the recipient of the message, "messageArg1:Arg2:" is the message to send to that object, and "var1", "var2", etc, are arguments to pass along with the message. Objective C's messages are keyword/value oriented, which is why the message name "messageArg1:Arg2:" is interspersed with the arguments.

The whole idea of Swarm is to provide an execution context within which a large number of objects can "live their lives" and interact with one another in a distributed, concurrent manner. Furthermore, we wish to insulate the user from having to master all of the highly baroque computer-science wizardry usually required to implement such massively distributed systems of autonomous agents reliably and robustly.

In the context of the Swarm simulation system, the generic outline of an experimental procedure takes the following form:

1. Create an artificial universe replete with space, time, and objects that can be located, within reason, to certain "points" in the overall structure of space and time within the universe., and allow these objects to determine their own behavior according to their own rules and internal state in concert with sampling the state of the world, usually only sparsely.
2. Create a number of objects which will serve to observe, record, and analyze data produced by the behavior of the objects in the artificial universe implemented in step 1)
3. Run the universe, moving both the simulation and observation objects forward in time under some explicit model of concurrency.

4. Interact with the experiment via the data produced by the instrumentation objects to perform a series of controlled experimental runs of the system.
5. Depending on what is observed in stage 4), alter the experimental or instrumental "apparatus" and go back to 3).
6. Publish paper *including* detailed specification of the experimental set-up so that others can recreate your experiment and verify your results.

Mag 3x: Swarm Structures

Swarm applications are structured around the concept of the *Swarm*. *Swarms* are the basic building blocks of Swarm simulations: a *Swarm* is a combination of a collection of objects and a schedule of activity over those objects. The collection are like the matter of the Swarm and the schedule is like the arrow of time moving the objects forward.

Model Swarms

For our current demos, Swarm applications contain two swarms. At the core is the *model swarm*, the Swarm that encapsulates the simulated model. Everything in the model swarm corresponds to objects in the world being modeled. For instance, in Heatbugs the HeatbugModelSwarm contains a collection of Heatbug agents, a HeatSpace to represent a physical property of the world, and a Grid2d to store agent position.

In addition to the object collection, the model swarm also contains a schedule of activity on the model. The schedule defines the effect of passing time on the model. For the simple heatbugs schedule, the execution is simply to update the HeatSpace (diffusing heat across the world) and then telling each Heatbug agent to move itself.

Model swarms consist of a set of inputs and outputs. The inputs to the HeatbugModelSwarm are model parameters: things like the size of the world, the number of HeatBugs, and the diffusion rate of heat. The outputs of the HeatbugModelSwarm are the observables of the model: the individual Heatbugs, the distribution of heat across the world, etc.

ObserverSwarms

The model swarm alone defines the simulated world. But an experiment does not just consist of the objects being experimented upon, it also includes the experimental apparatus used for observation and measurements. In Swarm computer simulations, those observation objects are placed in an *observer swarm*.

The most important object in an observer swarm is the model swarm that is being studied. The model swarm is one component of the observer, kind of like a little world in a petri dish on the lab bench. Other observer objects can then input data into the model swarm (setting simulation parameters, for instance) and read data out of the model swarm (collecting statistics of the behavior of agents).

Just as in setting up a model swarm, an observer swarm has a collection of objects (the instrumentation), a schedule of activity, and a set of inputs and outputs. The activity of the observer schedule is to drive data collection - read this number out of the model, draw it on a graph. The inputs to the observer swarm are configurations of the observer tools: what sorts of graphs to generate, for instance. The outputs are the observations.

When running in graphics mode, the observer swarm objects are largely used to mediate user interface. For instance, in Heatbugs the HeatbugObserverSwarm creates Raster widgets, EZGraphs, and Probes. All of these objects are connected into the HeatbugModel swarm to read data, and to graphical interface objects so the human sitting in front of the computer can observe the world.

Interactive, graphical experimentation with models is useful for coming up with intuitions. But for serious experimentation it is necessary to collect statistics, which means doing many runs and storing data for analysis. As an alternative to a graphical observer swarm, you can also create *batch swarms*, observer swarms that are intended to be run without any interaction at all. Instead of putting up fancy graphics, batch swarms read data from files to control the model and write the data out to other files for analysis. The key here is that the model swarm used in the batch swarm is the exact same model as that used in a graphical observer swarm: the only difference is what tools the observer (batch or graphical) connects to the model.

Summary

Multiple Swarms are used to create an experimental apparatus and control it. The use of multiple Swarms is not restricted only to this use, though: in particular, a model Swarm could itself contain its own subswarms, building a hierarchical simulation. In future Swarm development, we intend to use the power of the multiple Swarm modeling approach to build complicated and flexible models.

Mag 4x: Sketch of Code

Now that we have multiple Swarms and an experimental apparatus, it's time to learn how to use the objects themselves inside an application. Some examples are provided here: to understand this better, it will be necessary to read through example applications and the library documentation. These examples come from the Heatbugs application.

Building a Model Swarm

The key component of a simulation is the model Swarm. Here is the definition of a HeatbugModelSwarm, from HeatbugModelSwarm.h

```
@interface HeatbugModelSwarm : Swarm {
 int numBugs; // simulation parameters
 double evaporationRate;
 double diffuseConstant;
 int worldXSize, worldYSize;
 int minIdealTemp, maxIdealTemp;
 int minOutputHeat, maxOutputHeat;
 double randomMoveProbability;

 id modelActions; // scheduling data structures
 id modelSchedule;

 id heatbugList; // list of all the heatbugs
 Grid2d * world; // objects representing
 HeatSpace * heat; // the world
}

-getHeatbugList; // access methods into the
-(Grid2d *) getWorld; // model swarm. These methods
-(HeatSpace *) getHeat; // allow the model swarm to
 // be observed.

+createBegin: (id) aZone; // extra methods you
-createEnd; // provide for Swarms
-buildObjects;
-buildActions;
-activateIn: (id) swarmContext;
```

The first section of code says that a HeatbugModelSwarm is a kind of Swarm. HeatbugModelSwarm inherits a lot of behavior from generic Swarm, but also adds new variables and methods.

The new variables are enclosed in the braces in the definition of HeatbugModelSwarm. They are split into three general classes of things: simulation parameters, schedule data structures, and objects in the world. This is a typical sort of model swarm.

Finally, a HeatbugModelSwarm defines new methods. The first few methods are used to allow the model to be observed: a HeatbugModelSwarm will give out its list of Heatbugs, for instance, or its HeatSpace. Observers use these methods to monitor the model.

In addition to the observation methods, there are several Swarm-specific methods for the building of Swarms. These are fairly stereotyped. The `createBegin` and `createEnd` messages are used to create the Swarm object itself. `buildObjects` builds the model objects, and `buildActions` builds the model schedule - more on these later. Finally, `activateIn` arranges for the execution machinery to execute the Swarm itself.

Defining an Agent

The agents are typically the real focus of a modeling effort. Most of the work in a simulation comes in defining the agent behavior so that the computer agents resemble the real world phenomena you are trying to create.

In the case of Heatbugs, agents are pretty simple. Here is their definition, from Heatbug.h:

```
@interface Heatbug: SwarmObject {
 double unhappiness; // my current unhappiness
 int x, y; // my spatial coordinates
 HeatValue idealTemperature; // my ideal temperature
 HeatValue outputHeat; // how much heat I put out
 float randomMoveProbability; // chance of moving randomly

 Grid2d * world; // the world I live in
 int worldXSize, worldYSize; // how big that world is
 HeatSpace * heat; // the heat for the world
 Color bugColor; // my colour (display)
}

-setWorld: (Grid2d *) w Heat: (HeatSpace *) h; // which world are we in?
-createEnd;

-(double) getUnhappiness;

-setIdealTemperature: (HeatValue) i;
-setOutputHeat: (HeatValue) o;
-setRandomMoveProbability: (float) p;
-setX: (int) x Y: (int) y; // bug's position
-setBugColor: (Color) c; // bug's colour (display)

-step;

-drawSelfOn: (Raster *) r;
```

Heatbug is a subclass of SwarmObject. SwarmObjects have very little behavior of their own - they are defined as the root class of most objects and control computer science aspects like memory allocation and probability.

Heatbug carry with them a variety of state variables. For instance, each Heatbug has a notion of its ideal temperature, which will affect it's behavior. In addition, Heatbugs have variables that let them know about the world: these agents are storing references to the HeatSpace object, for example.

Most of the Heatbug methods have to do with setting up the agents state - the inputs to a Heatbug. Every heatbug must set up its world and heat objects, via the `setWorld:Heat:` method. In addition when Heatbugs are created they have their ideal temperature set, their output heat, etc. Heatbugs are also observable. Heatbugs define a getUnhappiness method - the unhappiness is the major measurable aspect of a heatbug, how well optimized it is at the moment. They also have a drawSelfOn method that directs the heatbug to draw itself on the specified graphics widget.

Finally, and most importantly, a Heatbug has a step method. step is where the Heatbugs behavior is defined: every time the Heatbug is told to step it performs its internal calculations, choosing where to move. Each heatbug is told to step when appropriate by the model schedule. The code for step is the real intellectual input into the model, and is worth reading as an example of an agent's behavior.

Building Agents

Now that Heatbugs have been defined, the model swarm needs to create them. This code fragment is from the buildObjects method on HeatbugModelSwarm.

```
// A loop to create a bunch of heatbugs.
for (i = 0; i < numBugs; i++) {
 Heatbug * hbug;
 int idealTemp, outputHeat;

 // Choose a random ideal temperature, output heat from the specified
 // range (model parameters).
 idealTemp = [uniformRandom rMin: minIdealTemp Max: maxIdealTemp];
 outputHeat = [uniformRandom rMin: minOutputHeat Max: maxOutputHeat];

 // Create the heatbug, set the creation time variables
 hbug = [Heatbug createBegin: [self getZone]];
 [hbug setWorld: world Heat: heat];
 hbug = [hbug createEnd];

 // Add the bug to the end of the list.
 [heatbugList addLast: hbug];

 // Now initialize the rest of the heatbug's state.
 [hbug setIdealTemperature: idealTemp];
 [hbug setOutputHeat: outputHeat];
 [hbug setX: [uniformRandom rMax: worldXSize] // random position
 Y: [uniformRandom rMax: worldYSize]];
}
```

The details of this code are best explained in reading the documentation for the libraries. Essentially, we first generate two random numbers: an ideal temperature and an output heat for the new Heatbug. We then create the Hheatbug itself with createBegin and fill in the required parameters of world and heat. Once those are set, we can send createEnd to the Heatbug and it is finished being created. After it's done being created we add it into a list of Heatbugs in the model and set a few parameters on it like the ideal temperature and the initial position.

Building Space objects

In Swarm, spaces are really just another kind of agent. In the heatbugs model we create a HeatSpace, a subclass of a diffusion object from the Swarm space libaries (specified in HeatSpace.m). Here is the code from `buildObjects` in the HeatbugModelSwarm

```
heat = [HeatSpace createBegin: [self getZone]];
[heat setSizeX: worldXSize Y: worldYSize];
[heat setDiffusionConstant: diffuseConstant];
[heat setEvaporationRate: evaporationRate];
heat = [heat createEnd];
```

the object is created, a few parameters are set, and then the creation is finalized.

Scheduling a Model Swarm

Once all of the simulated objects are created in `buildObjects`, the next task is to schedule them in the method `buildActions`.

```
modelActions = [ActionGroup create: [self getZone]];
[modelActions createActionTo: heat message: M(stepRule)];
[modelActions createActionForEach: heatbugList message: M(step)];
[modelActions createActionTo: heat message: M(updateLattice)];

modelSchedule = [Schedule createBegin: [self getZone]];
[modelSchedule setRepeatInterval: 1];
modelSchedule = [modelSchedule createEnd];
[modelSchedule at: 0 createAction: modelActions];
```

The heatbug model schedule actually consists of two components: an ActionGroup called `modelActions` and a Schedule called `modelSchedule`. The ActionGroup is a tightly coupled list of three messages: every time the action group is executed, it will send three messages in a row:

```
[heat stepRule];
[heatbugList forEach: step];
[heat updateLattice];
```

The ActionGroup alone specifies three messages to send - in order to put it in the simulation, that ActionGroup is then dropped into a Schedule. The Schedule itself only has one action - to execute `modelActions` itself. That action takes place at time 0. But because we've set a repeat interval on the schedule of 1, the schedule itself loops, executing every 1 time step. The final result is that `modelActions` is executed at time 0, time 1, etc.

Building a Graphical Observer Swarm

With the model swarm defined, arranging for a graphical observer Swarm is the next step. For Heatbugs, the code is in HeatbugObserverSwarm. The structure of an observer swarm is almost exactly like building a model swarm.

```
@interface HeatbugObserverSwarm : GUISwarm {
 int displayFrequency; // one parameter: update freq

 id displayActions; // schedule data structs
 id displaySchedule;

 HeatbugModelSwarm * heatbugModelSwarm; // the Swarm we're observing

 // Lots of display objects. First, widgets
 XColormap * colormap; // allocate colours
 ZoomRaster * worldRaster; // 2d display widget
 EZGraph * unhappyGraph; // graphing widget

 // Now, higher order display and data objects
 Value2dDisplay * heatDisplay; // display the heat
 Object2dDisplay * heatbugDisplay; // display the heatbugs
}
```

Again we have input parameters (display frequency), schedule data structures, and resident objects (model swarm, display widgets). The important exception is that HeatbugObserverSwarm is a subclass not just of the generic Swarm class, but specifically a GUISwarm. That implies that the HeatbugObserverSwarm will contain a control panel to allow the user to stop execution, and will also have a special `go` method to set everything running.

Building a Data Graph

An example of an object inside the HeatbugObserverSwarm is a data graph, the graph of average unhappiness. Here is the code necessary to create that object:

```
// Create the graph widget to display unhappiness.
unhappyGraph = [EZGraph createBegin: [self getZone]];
[unhappyGraph setTitle: "Unhappiness of bugs vs. time"];
[unhappyGraph setAxisLabelsX: "time" Y: "unhappiness"];
unhappyGraph = [unhappyGraph createEnd] ;

[unhappyGraph createAverageSequence: "unhappiness"
 withFeedFrom: [heatbugModelSwarm getHeatbugList]
 andSelector: M(getUnhappiness)] ;
```

The first step is to build an instance of an EZGraph and set its captions. Then a Sequence is created inside that graph (in this case, an AverageSequence). In general, the Sequence requires a target object and a message with which to extract data from that object -- the data is then plotted as one line in the graph.

> In the case of an AverageSequence, an entire collection of objects is presented to it. The AverageSequence then extracts data from all the objects in the collection

(in this case a List) using the provided message (in this case getUnhappiness), and generate a datapoint from these values by averaging them.

main()

The last main type of code needed for an application is the function main(), the first function called in your program. All the real work has been done already - all that's left is to create the objects at the right time.

```
int
main(int argc, char ** argv) {
 id theTopLevelSwarm ;

 // Swarm initialization: all Swarm apps must call this first.
 initSwarm(argc, argv);

 // swarmGUIMode is set in initSwarm(). It's set to be 1 if you
 // typed heatbugs -batchmode. Otherwise, it's set to 0.

 if (swarmGUIMode == 1)
 // We've got graphics, so make a full ObserverSwarm to get GUI objects
 theTopLevelSwarm = [HeatbugObserverSwarm create: globalZone];
 else
 // No graphics - make a batchmode swarm and run it.
 theTopLevelSwarm = [HeatbugBatchSwarm create: globalZone];

 [theTopLevelSwarm buildObjects];
 [theTopLevelSwarm buildActions];
 [theTopLevelSwarm activateIn: nil];
 [theTopLevelSwarm go];

 // theTopLevelSwarm has finished processing, so it's time to quit.
 return 0;
}
```

main() calls initSwarm (required in all Swarm applications). It then detects if it should do graphics or not, creates the appropriate top level Swarm to contain the model, and sets it to running. Simple as that!

Conclusion

Swarm tries to help computer simulation authors by making it easier to write simulations by making it more formal. The text above gives a narrative introduction into using Swarm for your own models, but a real understanding of Swarm will only come when you start to read through our examples and try to write your own applications. Good luck, and may your simulations be successful!

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Installing Swarm

Installing Swarm will take some time: various external libraries need to be acquired, compiled, and installed, and then Swarm itself needs to be compiled. Please report any problems during installation back so that we can try to fix them in the future. Also we've tried to compile some of the common obstacles to installation in the Swarm FAQ

(faq/FAQ.html)
.

- Supported Systems
- Prerequisite Programs
- Prerequisite Libraries
- Configuring Swarm Libraries
- Compiling Swarm Libraries
- Compiling and Running Swarm Applications
- Post-compile cleanup

Supported Systems

Our intention is that Swarm will run on any modern Unix system. Ideally, Swarm itself should be 100% portable. Swarm is currently known to be running on SunOS 4.1.3, Linux, HP/UX, some SGI Irix boxes, MachTen/68k, and Solaris. Some of the software Swarm uses might not be entirely portable; these problems are largely out of our hands. The major difficulty so far is the tricks the GNU Objective C runtime uses may not work properly on AIX and MachTen/PPC, so Swarm may not work on those systems. Fixing this is a longterm project; if you absolutely must use Swarm on AIX or MachTen/PPC, contact us.

> See the Platform News
> (http://www.santafe.edu/projects/swarm/platforms.html)
> web page for more dynamic information.

Prerequisite Programs

We assume you already have Unix and X Windows on your system: the rest of this software is freely available. Expect to spend some time compiling and installing these packages. URLs on this page are to the home distribution site: as a convenience, the Swarm ftp site ftp://ftp.santafe.edu/pub/swarm/

(ftp://ftp.santafe.edu/pub/swarm/needed-software/)
contains copies of the necessary packages in needed-software/
(ftp://ftp.santafe.edu/pub/swarm/needed-software/)
.

gcc 2.7.x

gcc is the FSF's C, C++, and Objective C compiler. You need gcc 2.7.2 or greater installed on your system, including the Objective C support, libobjc.a. The latest version is available from ftp://prep.ai.mit.edu/pub/gnu/ (ftp://prep.ai.mit.edu/pub/gnu) . Note: Swarm will not work with versions of gcc before 2.7.2.

GNU make

Make is used to automate building programs: every OS ships their own version of make with some random set of features. The Swarm makefiles currently use features that are not supported by all makes: GNU make is the only one guaranteed to work. Linux systems probably use GNU make already. The latest version of GNU make should be available on ftp://prep.ai.mit.edu/pub/gnu/ (ftp://prep.ai.mit.edu/pub/gnu/)

gdb

gdb is the GNU debugger: Swarm doesn't require this, but you will probably want gdb on your system to debug programs (debugging.html) . Gdb is available on ftp://prep.ai.mit.edu/pub/gnu/ (ftp://prep.ai.mit.edu/pub/gnu/) .

Prerequisite Libraries

Swarm uses a variety of freely available software libraries. All of these libraries need to be configured and installed independently of Swarm. When you have installed them all, modify Makefile.conf in the Swarm library package to tell the compiler and linker where to find them. Note: it pays to spend the extra time to build shared library versions of these things.

libXpm

libXpm adds pixmap (coloured bitmap) support to X11. libXpm is a common X library, many systems already have it installed. A reasonably modern version is needed: we've used version 3.4f. If you get complaints about "Object" being multiply defined, your version of libXpm is too old. The library is available from ftp://ftp.x.org/contrib/libraries/xpm-3.4h.tar.gz (ftp://ftp.x.org/contrib/libraries/xpm-3.4h.tar.gz) .

Tcl/Tk

More common libraries: a scripting language and a GUI toolkit. These libraries are available from ftp://ftp.smli.com/pub/tcl/ (ftp://ftp.smli.com/pub/tcl/) or from http://www.tcltk.com (http://www.tcltk.com)

. Swarm works with Tcl/Tk versions 7.4/4.0, 7.5/4.1, and 7.6/4.2. It may work with later versions, but it hasn't been tested.

BLT-2.3

BLT is an add-on to Tk that provides more widgets. BLT is available from http://www.tcltk.com/blt/index.html (http://www.tcltk.com/blt/index.html) or ftp://ftp.neosoft.com/pub/tcl/sorted/devel/ (ftp://ftp.neosoft.com/pub/tcl/sorted/devel/) . Note that BLT 2.3 works with either Tcl/Tk 7.5/4.1 or 7.6/4.2, whereas BLT 2.1 only works with versions up to and including Tcl/Tk 7.5/4.1. There are problems with using BLT 2.3 in that some of the functionality has changed and non-fatal errors may spit out on the terminal window you start Swarm in. But, we haven't taken the time to debug these things.

tclobjc-1.2

Glue between TCL and Objective C. This version of the library is currently only available from the Swarm authors, and is distributed (http://www.santafe.edu/projects/swarm/release.html) in the same place as Swarm.

Configuring Swarm Libraries

Once you have all of the above software installed, the next step is to configure Swarm. First unpack the Swarm .tar.gz file into a convenient place (via gzip -dc swarm-xx.yy.zz.tar.gz | tar xf -).

All configuration of Swarm compilation is localized to the file Makefile.conf. You should change the lines OTHERINCDIRS and OTHERLIBDIRS to point to the directories where the libraries you installed are, and set OTHERLIBS to link in other OS-dependent libraries that might be required on your system. In addition, if your OS doesn't have ranlib, then configure the RANLIB= line to just execute /bin/true.

All of the procedures run by make are defined in the three top level makefiles Makefile.conf, Makefile.lib, and Makefile.appl. If you run into problems, you might need to look at the Makefile code in there.

Compiling Swarm Libraries – (not required for binary distributions)

You've done the hard part, now type "make" and watch the fun! Make will recursively run make in subdirectories, building libraries, installing header files in include and libraries in lib. If running "make" in the root Swarm directory does nothing, you probably aren't running GNU Make.

Swarm is set up so that you can run make in subdirectories and expect it to work. If a particular directory doesn't seem to make correctly, try to cd into that directory and run "make" from there to better see what the problem is.

Compiling and Running Swarm Applications

Swarm applications are distributed separately: you will need to download and unpack applications independently. After the application is unpacked, you will have to change the SWARMHOME variable in the Makefile. It is setup, by default to point to "../swarm", as if the Swarm libraries were in a sibling directory called "swarm". Just edit that string to point to where you installed the Swarm libraries. No other configuration should be necessary. All you need to do to compile an application is type "make", and it will compile and link.

Once the application is compiled, just run it out of its own directory. A control panel and a couple of parameter windows should pop up; press the "Go" button to watch it run. Congratulations!

Post-compile cleanup

After you've tried out the Swarm libraries for awhile you may want to clean up the intermediate .o files. Swarm is set up so that the only files that you need after compile are `$(SWARMHOME)/Makefile.appl`, `$(SWARMHOME)/Makefile.lib`, `$(SWARMHOME)/Makefile.conf`, and the contents of `$(SWARMHOME)/lib` and `$(SWARMHOME)/include`. If you run `make clean` in `$(SWARMHOME)`, it will throw away all the .o files in the source directories. You can even copy out the Makefiles, `lib`, and `include` to someplace else on your system: no paths are compiled in. Just remember to set `$(SWARMHOME)` to the right directory in your application's makefile. If you need to recompile the library and want to start from a completely clean source directory, `make realclean` will delete the includes and libs created or copied during the compile process. Note that `make realclean` will delete *all* the files in the include and lib directories. So, if you've set your Makefile.conf to point to a directory where you have other include files and libraries, then you don't want to do a `make realclean`.

For debugging purposes it's easiest if your Swarm libraries are compiled with `OPTFLAGS=-g` and you leave the sources around. However, your applications will run faster if you compile with optimization, `OPTFLAGS=-O3 -g`. This will make debugging inside the Swarm libraries strange, but gives noticeable performance improvement. Users of Sparc based machines should also see the notes about `-mv8` in `Makefile.conf`.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Overview of the Swarm Distribution

Swarm is currently in beta test release. Installation directions can be found in install.html

(install.html)
. Documentation is incomplete: learning to program with Swarm will require reading example applications and looking at the header files in libraries. All existing documentation is indexed in index.html
(index.html)
.

Libraries

The main part of Swarm is a set of libraries, one per source directory. Briefly, the directory organization is as follows:

lib, include, bin

Installed libraries, include files, and helper binaries.

src

Swarm library sources.

defobj
(src/defobj.html)
: Support of Swarm style Objective C programming

collections
(src/collections.html)
: Collection library -- various objects that "collect" other objects

activity
(src/activity.html)
: Schedule execution mechanism -- the "process"

objectbase
(src/objectbase.html)
: Base class for simulated objects -- the "agents"

space
(src/space.html)
: 2d spatial environment libary

random
(src/random.html)

Random number library

simtools
(src/simtools.html)
Collected bits of code necessary for running Swarm

tkobjc
(src/tkobjc.html)
Graphical Interface library

analysis
(src/analysis.html)
Various objects that help with data processing

Applications

In addition, there is a set of sample Swarm applications that are distributed separately. Applications are the best current roadmap for the Swarm code: much of what is possible with Swarm is demonstrated in the sample applications. The most important applications are:

template
Template simulation. The code itself here is trivial, but provides a nice base for new Swarm programmers to start from.

heatbugs
Heatbugs, a simple complex system. The code here is thoroughly commented for use as a guide to swarm programming.

mousetrap
Mousetrap is a discrete event simulation of a room full of mousetraps loaded with ping-pong balls. The triggering of one of the traps sets off a chain reaction supposedly similar to fission. It is also thoroughly commented.

market
Market is an Objective-C wrapped piece of legacy code originally written in C. It is a useful example of how to convert legacy code into Swarm; but, it is not very instructive on Swarm programming in general.

Credits

The Swarm developers are Roger Burkhart <roger@ci.deere.com>, Glen Ropella <gepr@santafe.edu>, Marcus Daniels <mgd@santafe.edu>, and Alex Lancaster <alex@santafe.edu> under the direction of Chris Langton <cgl@santafe.edu>. But credit for the package is spread out over many people. The Swarm developers can be reached collectively as <swarm@santafe.edu>. More information is available on our Web page, http://www.santafe.edu/projects/swarm/

(http://www.santafe.edu/projects/swarm/)
.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Swarm Todo/Wish list

Swarm is an open ended system which is meant to grow in response to the requirements of the user base, either by inhouse development or through user re-contributions. We are therefore maintaining a list of the most popular requests (both in terms of tools and libraries) so that groups of users can recognize common requirements, make more informed suggestions and so forth:

A 'Double' Space

Which could deal with notions of "distance" and answer questions of the form: "which other objects are within X radius of me"? An initial implementation of such a space has been re-contributed by Ginger Booth (http://peaplant.biology.yale.edu:8001/~ginger/) and may serve as a good foundation for such a space.

Complete Batch-Mode Support

Swarm can now run in batch mode, which should allow the user to organize large parameter sweeps over the models s/he is implementing. However, we still need to provide adequate and standardized support for file operations (it should be easy, for example, to load the parameters of an experiment from a file). In order to do this we will provide File objects which will allow users to avoid ad-hoc coding of their file-I/O. This support will be in place well before V1.0.

More Analysis Tools

The averager object can generate the Mean, Max, Min, and Count of a given input stream. We would like to add similar tools to calculate entropies, mutual information and other such measures.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Debugging Tips for Swarm

Debugging software is an art, perhaps even more so than writing the software itself. All software we write will have bugs: it is important to know how to diagnose a bug when it happens and how to write code defensively so you have less bugs in the first place.

Finding bugs

Bugs come in two categories: those that crash your program and those that don't. Bugs that crash your program are friendlier because they're obvious. Bugs that you don't notice are much more dangerous: one nagging question in every programmer's mind should be "is the program doing what I think it is doing?"

gdb

By far the most useful tool for finding bugs is a good debugger, a shell you can run a program under and set breakpoints, inspect the values of variables, etc. The best free debugger for Unix is probably gdb, available from the GNU ftp site

(ftp://prep.ai.mit.edu/pub/gnu/)
. gdb seems unfriendly and confusing at first, but it is definitely worth your time to learn it.

> The most important gdb commands are `help`: browse online help, `where`: show me a stack trace, where did we crash?, `list`: show me the source code where we are, `break`, set a breakpoint, and `print`: show me the value of of some expression. If your program is crashing on you, run it under gdb and look at the stack trace. If it looks to be buggy but you don't know where, start setting breakpoints and see when things go awry.

gdb and Objective C

Unfortunately, at this time gdb does not directly support Objective C. There are some workarounds

(http://www.cons.org/cracauer/objc-hint-gdb.html)
that make debugging Objective C programs possible. They are based on the knowledge that Objective C is little more than a glorified syntax for structs (objects) and strange function names (methods).

defobj: xprint(), xprintid(), xfprint(), xfprintid(), xexec(), and xfexec()

Swarm also has a few functions defined that can be used to make debugging easier. In particular, the function xprint(object) prints out the class of an object, and xexec(object, "message") calls the message specified on the object. These can be invoked under gdb as `call xprint(aHeatbug)`.

Note that you can't pass arguments to a message, nor can you see the return value. There are also methods xfprint(collection) and xfexec(collection, "message") that print or exec foreach member of a collection.

Preventing Bugs

Defensive programming can help prevent a good number of bugs. When writing code, try to test it incrementally: make small changes whose effects you think you can predict and then test them. Don't outsmart yourself with cleverness: write code correctly first, then go in and hack it up if you need it to be more efficient. Put in sanity checks for conditions that shouldn't go wrong in normal usage, but might if you make a mistake.

-Wall

Swarm currently compiles all code with "gcc -Wall", which tells gcc to emit warnings for a lot of things that it wouldn't normally complain about. Warnings are not (necessarily) errors - warnings will be generated for legal code if gcc thinks that what you're doing could easily be a mistake. You might find this frustrating at first, but it helps catch a lot of common errors, including forgetting to include a prototype or forgetting to return a value from a function. Passing -Wall is good discipline.

nil˙method

Objects in Objective C are essentially pointers to structs. So what happens if you send an object to the pointer 0x0, "nil" in Objective C parlance? Unfortunately, most implementations of Objective C, including gcc, define methods to nil as having no effect. The code

```
aHeatbug = [Heatbug create: aZone];
 aHeatbug = 0; // oops! bug.
 [aHeatbug setIdealTemperature: idealTemp];
```

will not generate any errors.

The reason this is unfortunate is that it's a common bug to trash a pointer accidentally, set it to 0. It would be nice if your program then crashed when you tried to send a message to that mangled object: instead, the message send will fail silently and the program will continue to execute. This can make it hard to find bugs.

There are two ways to make messages sent to nil crash your program. The simplest is to put a breakpoint on `nil_method` under gdb: nil_method is invoked every time a message is sent to nil. Alternately, you can make a copy of the libobjc runtime source and edit nil_method to do whatever you want. The source code in `src/defobj` contains a file `objc.patch` that patches the runtime from gcc 2.7.2.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Emacs and Objective C

Emacs is the world's best editor, infinitely configurable and powerful and free. If you're not used to a Unix text editor and want to write programs, emacs is probably best to learn. The information here will show you how to set up emacs to use an Objective C specific mode for editing source code and how to get emacs to colour your source code for you nicely. Emacs can also help you find compilation errors, run a debugger, and even act like a class browser. Docs for those things are not here, see the Emacs documentation itself.

objc-mode

One of the best things about emacs is that one can load different "major modes" to edit types of files. Major modes making editing files easier by providing structure to your editing. In particular, there are programming language modes that do nice things like indentation for you.

> Starting with emacs-19.30, the default C mode is cc-mode.el, a nice rewrite of the original c-mode.el. Luckily for us, cc-mode.el supports Objective C directly. The basic function to invoke it is `objc-mode`. The following bit of code in your .emacs will cause emacs to automatically enter objc-mode on all files that end in `.m` or .h:

```
(setq auto-mode-alist
      (append '(("\\.h$" . objc-mode)
                ("\\.m$" . objc-mode))))
```

hilit19

Emacs running under X has the ability to colour your source code according to syntax. There are two packages to do this: font-lock and hilit19. I'm told font-lock is better, but I still use hilit19. Once you have hilit19 running in your emacs, you need to teach it how to highlight Objective C code. Here's what I use (cut and paste)

```
;; modified objective c - get defuns, id.
  (let ((comments     '(("/\\*" "\\*/" comment)))
(c++-comments '(("//.*$" nil comment)
("^/.*$" nil comment)))
(strings      '((hilit-string-find ?' string)))
(preprocessor '(("^#[ \t]*\\(undef\\|define\\).*$" "[^\\]$" define)
("^#.*$" nil include))))

    (hilit-set-mode-patterns
     '(objc-mode objective-C-mode)
     (append
      comments c++-comments strings preprocessor
```

```
      '(
;; function decls are expected to have types on the previous line
("^\\(\\(\\w\\|[$_]\\)+::\\)?\\(\\w\\|[$_]\\)+\\s *\\(\\(\\w\\|[$_]\\)+\\s *((\\|(\\)[^)]*)+" nil defun)
("^\\(\\(\\w\\|[$_]\\)+[ \t]*::[ \t]*\\)?\\(\\(\\w\\|[$_]\\)+\\|operator.*\\)\\s *\\(\\(\\w\\|[$_]\\)+\\s *((\\|(\\)[^)]*)+" nil defun)
;; nelson - objective c style defuns
("^[-+].*{" nil defun)

("^\\(template\\|typedef\\|struct\\|union\\|class\\|enum\\|public\\|private\\|protected\\).*$" nil decl)
;; datatype -- black magic regular expression (nelson - add id)
("[ \n\t({]\\(\\(const\\|register\\|volatile\\|unsigned\\|extern\\|static\\)\\s +\\)*\\(\\(\\w\\|[$_]\\)+_t\\|float\\|double\\|void\\|char\\|short\\|int\\|unsigned\\|long\\|id\\|FILE\\|\\(\\(struct\\|union\\|enum\\|class\\)\\([ \t]+\\(\\w\\|[$_]\\)*\\)\\)\\)\\(\\s +\\*+)?\\|[ \n\t;()]\\)" nil type)

;; Objective C stuff: ^@[word] gets keyword colouring
("^@\\sw*" nil keyword)
;; key words
("[^_]\\<\\(return\\|goto\\|if\\|else\\|case\\|default\\|switch\\|break\\|continue\\|while\\|do\\|for\\|super\\|self\\)\\>[^_]"
 1 keyword)))))
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library

Documentation and Implementation Status

Overview

The defobj library supports the style of object-oriented programming that is used throughout Swarm. It defines a specific style for using the Objective C language that includes its own standard conventions for creating objects and for storage allocation, error handling, and debugging support.

Documentation

The defobj library follows the Swarm library interface conventions

(defobj/library.html)
. It has following standard sections of documentation:

- Usage Guide
(defobj/guide.html)
- Advanced Usage Guide
(defobj/adv.html)

- Interface Reference
(defobj/ref.html)
- Subclassing Reference
(defobj/subclass.html)

- Interface Design Notes
(defobj/design.html)
- Implementation Notes
(defobj/impl.html)

Marcus G. Daniels <mgd@santafe.edu>

(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Documentation and Implementation Status

The current interfaces defined in defobj are expected to remain stable, with the exception of those relating to package structure and custom-generated classes. Those facilities are in process of being totally replaced, and will not be documented further until that replacement is complete. Details of support for defining a customized type will also be provided later.

All the interfaces defined in the header file are currently implemented, except for the various types of zones that would hold all allocated storage in a collection of local pages, and that would provide automatic reclaim of unused storage. The current zone implementation supports all defined allocation messages, and also maintains the population of the zone that consists of all objects created directly within it.

The Usage Guide section is only a start, but it does include subsections that discuss the style of Objective C programming adopted for use throughout Swarm. There is a full set of Interface Reference sections, though some details remain to be filled in.

Throughout the documentation, a parenthesized comment that starts with (.. indicates an editorial comment on the current status of implementation or documentation.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Library Interface Conventions

Overview

Some of the basic Swarm libraries are implemented using object definition conventions established by the defobj library

(../defobj.html)
. This document explains how to read and interpret the header files and interface definitions published by such libraries. It also explains the typical structure of documentation provided for libraries that follow these conventions.

These conventions are documented under the following subsections:

- library header file
- object type definitions
- GridTurtle example
- global object symbols
- interface design conventions
- documentation structure

Library Header File

For a library that adopts these conventions, a library is not merely a collection of source files that are compiled into a library archive under control of a make file. Instead, Objective C source files in the library are processed in a special way to publish their definitions not only as header files, but as generated objects available at runtime to make full use of the library. A library processed in this way is referred to as as a "package." (Note: some of the newer Swarm libraries, such as objectbase and random, don't undergo this special processing and yet still follow all the library interface conventions described in this document.)

The defobj
(../defobj.html)
library documents the full details of special processing performed on a package. For simply using a library, the key fact to keep in mind is that the entire public interface to a library is declared in the one header file having the same name as the package itself, plus a trailing .h suffix. Additionally, the header file of a library normally documents only its public interface, in a way that is completely separated from the implementation of the objects it specifies.

The separation of implementation means that a library publishes its interface entirely without reference to any Objective C classes which implement its objects. Even though classes are often thought of as separating the interface of an object from its implementation, this separation is far from complete. Not only do classes typically contain many internal methods not intended for external use, but they also define a particular storage format for an object defined as instance variables.

A library instead publishes its interface as a set of public object types. These object types may also be supplemented by global object constants called symbols. Both these kinds of definitions normally appear only in the header file of a library. The remaining source files in a library normally contain the classes which implement the object types.

In a library, the files which implement classes (including class header files), need not ever be referenced simply to make use of the implemented capabilities of the library object types. Documentation for the library is normally expressed entirely in terms of the types and symbols published in the library header file. If a feature does not appear in the library header file, it should not be considered part of a supported public interface.

Individual class header files are required to subclass from existing implementations, but interfaces for subclassing are an entirely separate issue from normal public use of an object library. Class inheritance can be a powerful implementation technique, but extension of an existing class framework is typically safe only if performed in explicitly permitted ways. If a library supports subclassing at all, it must carefully state which classes may be subclassed and in what ways. For a library package, this information is supplied outside the library header file. The library header file specifies only the interfaces by which objects are intended to be used, whether implemented by a local class or an external subclass.

Remaining sections of this document explain the declarations which appear in a library header file, and end with a suggested structure of documentation to be provided for a library. The libraries of Swarm mostly follow this structure.

Object Type Definitions

The interfaces to objects defined by a library are specified by object type definitions. An object type defines only a set of messages which may be sent to an object over various phases of its lifetime. An object type makes no commitment to the classes that might be used to implement the object.

Multiple classes may all implement the same messages belonging to a type. The independence of types and classes means that different classes can provide alternate implementations of the same object type. For example, a particular implementing class might be selected to optimize the implementation for a particular case.

Object types are similar to protocols defined by Objective C, and the declarations appearing in a library header file are a minor adaptation of Objective C protocol syntax. A key difference from protocols is that object types are published as real external objects that may be used at runtime to create instances of a type. A futher difference is that the object types of the defobj library are divided into separate interfaces that define distinct phases of an object's life cycle.

GridTurtle example

An example will help illustrate the features of type definitions supported by the defobj library. Throughout documentation of the basic Swarm libraries, a series of running examples will be based on a simple type of object belonging to sample Swarm simulations. This type of object is an agent that move around on a two-dimensional grid, always moving in a current direction that it maintains internally. This agent is like a "turtle" of the original Logo system, except that its position is constrained to discrete integer values of its X-Y coordinates, and its direction is always one of the four orthogonal directions north, east, south, or west.

Following is a complete library header file for a library which defines such an object, called GridTurtle:

```
/*
Name:         GridTurtle.h
Description:  object type for Swarm example programs
Library:      grid
*/

#import <defobj.h>

@deftype GridTurtle <Create, Drop, CREATABLE>CREATING
- (void) setXLocMax: (int)xLocMax;
- (void) setYLocMax: (int)yLocMax;
SETTING
- (void) setXLoc: (int)xLoc;
- (void) setYLoc: (int)yLoc;
- (void) setXLoc: xLoc setYLoc: yLoc;

- (void) setDirection: direction;
USING
- (int) getXLoc;
- (int) getYLoc;

- getDirection;

- (void) move: (int)distance;
- (void) turn: (int)angle;  // angle measured in units of pi/2 (90 deg.)
- (void) print;
@end
```

```
id <Symbol> North, East, South, West;

#import "grid.xt"
```

An object type is defined by an @deftype declaration. (Note: newer libraries, including objectbase and random, now follow the library interface conventions without using this special @deftype tag. Instead they use just an ordinary @protocol declaration, but otherwise they follow all the structure explained in this document.) The syntax of such declaration is identical to that of an Objective C @protocol definition, except for the entirely uppercase keywords (CREATABLE, CREATING, SETTING, USING) appearing in the GridTurtle example above. All these modifications of Objective C syntax are accomplished by simple preprocessor macros; no extensions to the language compiler are involved.

When this library header file is processed (by a special rule in a make file), an external object id with the name GridTurtle is automatically published. The name of a defined type becomes an ordinary object that accepts specific messages defined by the defobj library. The defobj library
(../defobj.html)
explains the details of such messages; the only purpose here is to explain the basic sections of a deftype declaration.

deftype declarations follow the syntax as Objective C protocols for inheriting messages from each other: a list of names enclosed in angle brackets (e.g., <Create, Drop, ...> above) gives the names of other declared types containing messages to be supported by the new type as well. (These types referenced here are defined by the imported file <defobj.h>.) Like protocols, full multiple inheritance of types is supported. The same messages may be inherited any number of times through any path with no different effect than if inherited or declared just once, so long as no conflicts occur in any of their argument or return types.

The CREATABLE tag appearing in the inherited type list above is a special type which defines no messages of its own, but merely marks the type as one which supports direct creation of instances of the type. Without this tag, the only role of a type is to define messages for inheritance by other types. With this tag, the global type object has a particular implementation that supports object creation using standard messages defined in defobj.

The declared messages of the type may be separated into sections marked by the special uppercase tags such as CREATING, SETTING, and USING above. (Currently, these are the only such tags which may occur.) These sections each define messages belonging to a particular defined "interface" of the object type, which are further combined into distinct "phases" of an object lifecycle supported by defobj messages. Further explanation of the interfaces and phases defined by this example are provided in the Usage Guide
(guide.html)
of the defobj library.

The following line at the end of the <grid.h> header file above is required as part of the standard conventions for coding a library header file: This line adds no information to the header file, but is required as part of the system by which definitions are automatically generated.

```
#import "grid.xt"
```

Global Object Symbols

The <grid.h> header file above also contains the declaration:

```
id <Symbol>  North, East, South, West;
```

Lines that declare global id variables of type Symbol, EventType, Warning, or Error (using the angle bracket syntax of id variables conforming to a protocol) are processed somewhat like deftype declarations in that they also produce global id variables initialized to support particular messages defined by defobj. These global variables, however, are not used to define or implement other message interfaces, but only to define certain fixed capabilities referenced through their global object names.

If declared as a Symbol, as in the case here, the generated objects have no particular behavior of their own (other than the character string of their name), but only serve to define unique global id constants which may be used as distinct named values in messages. In this example, the current direction of a GridTurtle object is represented by one of the symbol names North, East, South, or West. These values are like the enum constants of the C language, except that they are defined as full Objective C objects, and may be used with further restrictions. An EventType, or a Warning or Error, defines a further subtype of a Symbol constant with further specialized messages documented in defobj.

Interface Design Conventions

A variety of rules on naming and declaration of object types, symbols, and messages are followed by many of the Swarm libraries. These rules help establish a basic consistency on the library interfaces. Following is a list of such conventions that apply to a public library interface, some but not all of which are derived from standard Smalltalk or Objective C coding practice:

- Names of global object constants are capitalized. In the public interface, such names include types and symbols.
- Recapitalization separates words of a compound name (e.g., GridTurtle). Underscores are generally not used.
- Message names start with a lower-case character, and are named using verbs. Nouns that represent gettable or settable components of object state (e.g., Direction of a GridTurtle), are prefixed by get or set to indicate the action being performed.
- The Smalltalk convention that a message return the receiver of a message if there

is no other specific return value is generally *not* followed. If there is no specific return value needed from a message the return type is declared (void).

Documentation Structure

The following standard sections of documentation are suggested for libraries that follow the strict interface vs. implementation separation of library packages, each with purpose and typical contents as given:

- Usage Guide
  Tutorial introduction to a library. Focuses on the most common uses in an order reflecting the needs of a first-time user. Most explanation by means of progressively elaborated examples. Serves as a guided tour of major library capabilities. No attempt at reference-style completeness.
- Advanced Usage Guide
  Continues overview of all significant capabilities of a library, including those which might be needed by advanced users customizing or extending built-in capability. Provides examples of specialized uses extending beyond normal, basic usage, but still using built-in features of the library framework. Helps to simplify the Usage Guide by providing a location for overflow of more advanced features.
- Interface Reference
  Complete, concise summary of all features of a library. Fills the role a Unix "man page" in providing comprehensive definition of library services. Little or no concern to provide path of tutorial introduction. Does not contain extended examples, at most only fragments of examples which serve needs of specification.
- Subclassing Reference
  Documents the rules for writing new classes which subclass from classes that implement the library. Because a library might use complex combinations of classes to implement the range of behaviors defined by its types, the classes in a library are not automatically usable as superclasses of user-defined classes. Each library documents which classes are available for use as superclasses, and the specific rules that must be followed when subclassing from these classes.
- Interface Design Notes
  Explains why choices were made the way they were in the design of the public interface to a library, including other alternatives considered. Serves as a "Rationale" document for the documented interface. May include historical background deliberately left out of the preceding sections. May also include references to other work.
- Implementation Notes
  Explains the structure of classes by which the types of a library are implemented. Summarizes the status of implementation if still incomplete and lists items of possible future work. Provides overview and high-level structure of the implementation in whatever ways would best guide a reader of implementation source code. May

also discuss tradeoffs considered along with references to related or supporting work.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Usage Guide

Contents

Why Swarm uses Objective C

Swarm uses object-oriented programming not only because this is a good way to build general-purpose software libraries, but because the very concept of an object is at the base of how you build a model in Swarm. To build a simulation in Swarm, the first step to define the various kinds of objects that can inhabit some real or artificial world, and the second step is to define the kinds of events that can occur to these objects, including all the different ways that the objects can interact with each other.

The basic concept of an object is that it responds to external events, and that the only thing that really matters about an object are the ways it can be observed responding to these events. In an object-oriented programming system, an object is represented by some uniquely identified chunk of data, and the events are the dynamic operations that can be made to occur on these objects.

Different object-oriented programming systems define the operations that occur on objects in different ways. For example, in C++ the operations are called "member functions" and look much like an ordinary function call in the C language. In Objective C, an operation on an object is called a "message" and has a special syntax by which you call it, but a message also has arguments and behaves in many respects like a call to a function.

The key requirement for Swarm is to be able to make all these operations happen to any object at any time, whenever they're supposed to occur within some simulated world. It stores these operations inside its own data structures, and when you run a simulation in Swarm, the Swarm system itself traverses these data structures to make the necessary operations happen at the proper time.

Swarm needs to be able to make any kind of action happen to any kind of object at any time, and to do this it needs very general-purpose structures that hold some kind of representation of the actions themselves. That's a special requirement that not every object language provides. You already have enough work to define the kinds of objects that can exist inside your model, and Swarm doesn't want you to have to define an even larger number of objects for every kind of event that might occur to your objects. Instead, it wants the object language to provide it with a representation of the operations on objects that you've already defined. Swarm can then store these representations in its own data structures, and make the operations happen whenever it needs to.

In Swarm, the representation of events that happen on objects are just as fundamental to the model as the objects themselves. That's why Swarm is a "discrete-event" simulation system. But if the object system doesn't provide a general-purpose enough representation of events as well as objects, there would be a lot more work to do. The Objective C system, through its special data type called a "message selector," provides a representation of operations that is flexible enough to do this, but the C++ language does not. (In C++, the compiler requires more information about an operation than the event structures in Swarm would be able to provide.)

There's a host of additional reasons why Objective C has also been a good match for the requirements of the Swarm system. Like C++, the operations on objects can be compiled to a very efficient form (though Objective C does require a little more overhead than C++ to get the operation started). This efficiency can be very important for a simulation, since simulations can run for very long periods of time to explore all the behavior that might occur within their simulated worlds. Other languages, such as Lisp, Smalltalk, and Java, also have the "dynamic message dispatch" feature that would make general-purpose event structures possible, but they still carry significant added overhead compared to C.

Objective C is also a very simple extension to the C language. Basic knowledge of C is already widespread, and a few days are typically all that is needed to learn the few additions of Objective C. (Really learning the concepts of objects, however, can take much longer, no matter what object language you try to use.)

Objective C was also a good choice for Swarm because it has a high quality, freely distributable implementation in the GNU C compiler. One of the main concerns about Objective C is that it isn't nearly as widely known or used as other languages like C++, but at least the GNU C compiler assures it will be available on machines where Swarm needs to run. The OpenStep system of the NeXT corporation, and the parallel GNUStep project, also help assure that Objective C is a living language.

There are even more powerful aspects of Objective C that Swarm takes advantage of. Some of these are described in the Advanced Usage Guide
(adv.html)

of the defobj library. But many others are at the heart of how Swarm uses Objective C, and so are dealt with in the rest of this Usage Guide. The entire purpose of the defobj library is to define a standard style for the use of Objective C in Swarm. This style is backed up by a library of foundation classes that Swarm provides to support this style.

Swarm provides its own foundation classes even for such basic operations as creating an object, and other support that user classes ordinarily receive from a builtin Object superclass. A good understanding of the defobj library is essential for Objective C programming in Swarm. Like any programming language, Objective C requires learning not only the rules of the language itself, but also a standard library of initial capabilities that you build from. Objective C doesn't really have a single official standard library (or language definition either, for that matter), but the NextStep foundation libraries (and the NextStep language definition) come very close to this status.

For a variety of reasons explained here and in the Advanced Usage Guide
(adv.html)
, Swarm doesn't use a standard library based on the NextStep foundation interface. You can still learn the Objective C language from other available sources, but you have to be careful to sort out the language level from the standard object and library support they also discuss.

The index page
(../../index.html)
of the Swarm documentation provides a variety of links to other Objective C resources. In particular, it provides a link to a complete on-line reference for the Objective C language as defined by NextStep and implemented by the GNU C compiler. None of the Swarm documentation attempts to duplicate this coverage of the basic language; Swarm assumes that you learn Objective C from other available sources. The NextStep reference(*Object-Oriented Programming and the Objective C Language*
(http://www.next.com/Pubs/Documents/OPENSTEP/ObjectiveC/objctoc.htm)
) is the single best available source, and it can also be ordered as a hardcopy book from the NeXT publications page
(//www.next.com/Pubs/Documents/OPENSTEP)
. Don't try to program in Objective C without it.

Swarm style of Objective C Programming

No matter where you've learned to program in Objective C, don't try to use Objective C in Swarm without understanding the special ways in which Swarm adapts its use of the language and the language runtime system. Your place to find this information is right here in the defobj library. Defobj not only provides the most basic layer of foundation class libraries for Swarm, but also serves as the place where all the rules and guidelines for the use of Objective C in Swarm are gathered together.

Some of these rules and guidelines don't even require any specific software to implement them, but are just conventions and recommended style that anyone programming in Objective C can follow when their goals are similar to those of Swarm, and to a large extent often do. These kinds of conventions are the focus of this section; following sections discuss aspects of Swarm Objective C style that depend on specific software support.

Lots more to come... For now, see Library Interface Conventions
(library.html)
for some of this information, and Swarm tutorials for much of the rest.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Advanced Usage Guide

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Interface Reference

Dependencies

The defobj library is defined and documented using the library interface conventions

(library.html)
established by the defobj library. It imports the Object superclass and other standard type definitions of the GNU Objective C runtime system.

> The collections library must always be linked along with the defobj library. Even though the interface definitions of defobj do not depend directly on collections, the implementations of these libraries both require the presence of the other. Only the collections library should be initialized directly; initialization of the collections library automatically initializes the defobj library as well.

Contents

- module definition conventions
(module.html)
- create combination messages
(combo.html)
- global portability assumptions
(globalport.html)
- global definitions
(globaldef.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

DefinedObject – standard support for all objects

Synopsis

```
//
// DefinedObject -- object with defined type and implementation
//
@deftype DefinedObject
- getZone;
- (BOOL) respondsTo: (SEL)aSel;
- getClass;

- (ref_t) addRef: (notify_t)notifyFunction withArgument: (void *)arg;
- (void) removeRef: (ref_t)refVal;

- (int) compare: anObject;

- perform: (SEL)aSel;
- perform: (SEL)aSel with: anObject1;
- perform: (SEL)aSel with: anObject1 with: anObj2;
- perform: (SEL)aSel with: anObject1 with: anObj2 with: anObj3;

- (void) setDisplayName: (const char *)displayName;
- (const char *)getDisplayName;

- (void) describe: outputCharStream;
- (void) xprint;
- (void) xprintid;
@end
```

Description

```
DefinedObject is the top-level supertype for all objects that follow the object programming
conventions of the defobj library.  The messages defined by this type are the only messages
which should be assumed to be automatically available on objects that follow these conventions.
In particular, use of messages defined by the Object superclass of the GNU Objective C runtime
should not generally be assumed because future implementations of some objects might not give
continued access to them.
```

The DefinedObject type defines a minimum of standard messages, and leaves to other types the definition of message that might or might not apply in any general way to particular objects.

Inherited behavior

The DefinedObject type inherits no behavior from any other type. It is the top of the interface type hierarchy for objects that follow the defobj programming conventions. All objects that follow these conventions should inherit from DefinedObject, either directly or indirectly through other types.

Defined behavior

```
- getZone;
- (BOOL) respondsTo: (SEL)aSel;
- getClass;
```

The getZone message returns the zone in which the object was created. The respondsTo: message returns true if the object implements the message identified by the selector argument. To implement a message means only that some method will receive control if the message is sent to the object. (The method could still raise an error.) The respondsTo: message is implemented by direct lookup in a method dispatch table, so is just as fast as a normal message send. It provides a quick way to test whether the type of an object includes a particular message. getClass returns the class that implements the current behavior of an object.

Defined behavior

```
- (ref_t) addRef: (notify_t)notifyFunction withArgument: (void *)arg;
- (void) removeRef: (ref_t)refVal;
```

These messages add or remove an external reference to an object that is notified whenever the object is changed or relocated. (More details to follow.)

```
- (int) compare: anObject;
```

```
- perform: (SEL)aSel;
- perform: (SEL)aSel with: anObject1;
- perform: (SEL)aSel with: anObject1 with: anObj2;
- perform: (SEL)aSel with: anObject1 with: anObj2 with: anObj3;
```

These messages are have implementations in the Object superclass, excppt for the three-argument form of perform:. They have local implementations in defobj, however, just to make subclassing from Object unnecessary. They are currently the only messages defined in Object that are recommended for use in Swarm.

```
- (void) setDisplayName: (const char *)displayName;
- (const char *)getDisplayName;
```

These messages assign a character string as a name that identifies an object for display or debug purposes. A global function, xsetname(), is also available to set the name by calling the function from the debugger.

```
- (void) describe: outputCharStream;
- (void) xprint;
- (void) xprintid;
```

The describe: message prints a brief description of the object for debug purposes to the object passed as its argument. The object passed as the outputCharStream argument must accept a catC: message as defined in String and OutputStream in the collections library. Particular object types may generate object description strings with additional information beyond the built-in default, which is just to print the hex value of the object id pointer along with the name of its class, and the display name of the object, if any. xprint prints the standard describe string to standard output; it is also called by the xprint function that may be called from the debugger. xprintid prints just a one-line describe string, consisting of the built-in default.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Create – create a customized instance of a type

Synopsis

```
//
// Create -- create a customized instance of a type
//
@deftype Create <DefinedObject>
CREATING
+ create: aZone;
+ createBegin: aZone;
- createEnd;

+ customizeBegin: aZone;
- customizeEnd;
- customizeCopy: aZone;
@end
```

Description

The Create supertype defines standard messages that provide a general-purpose protocol for creating new objects. These messages may be used either to create a new instance of a type in one message, or to bracket a series of messages that customize available options for an object to be created. The separation of create-time specifications from later behavior of an object gives substantial flexibility to adapt a generic type to particular needs.

These create messages may be implemented either by a type that hides its implementing classes, or directly by a class that adopts these messages as a uniform interface for creating objects. If implemented directly by a class, then the class object serves as the type object in all message descriptions that follow. Otherwise, a type object might be implemented in a variety of ways that guarantee only that published messages on a type are accepted.

In addition to the create messages defined by Create, an object type may support any other messages of any other names or calling conventions. These messages define only a standard method for creating new objects that other types are free to inherit and implement in conformance with a uniform convention. Further conventions are established elsewhere (create combination messages
(combo.html)

) for standard ways in which create messages that combine several steps can be combined.

Inherited behavior

The Create type has no inherited behavior beyond the universal messages of the DefinedObject

(DefinedObject.html)
type.

Defined behavior

```
+ create: aZone;
```

The create: message creates a new instance of a type with default options. The zone argument specifies the source of storage for the new object. The receiving object of this message is a previously defined type object. The message is declared as a class message (with a + declaration tag) to indicate that the message is accepted only by the type object itself rather than an already created instance of the type (which a - declaration tag otherwise defines).

The create: message returns the new object just created.

Defined behavior

```
+ create: aZone;
```

The create: message creates a new instance of a type with default options. The zone argument specifies the source of storage for the new object. The receiving object of this message is a previously defined type object. The message is identified as a class message (a + declaration tag) to indicate that the message is accepted only by the type object itself rather than an already created instance of the type, which most messages of a type (those with a - declaration tag) otherwise define.

The create: message returns the new object just created. This object is an instance of some class selected to implement the type. The class which a type selects to implement an object may be obtained by the getClass message, but is not otherwise visible to the calling program. A caller never refers to any class name when creating objects using these messages, only to type names, which are automatically published as global constants from any @deftype declaration.

```
+ createBegin: aZone;
- createEnd;
```

The createBegin: and createEnd messages bracket a series of messages that specify options for an object being created. The intermediate messages can set values defined as parameters of the type, or provide other forms of specification using available messages. A particular object type defines the specific messages that are valid for sending during this interim creation phase.

createBegin: returns an interim object intended only for receiving create-time messages. If a type was defined by a @deftype declaration, these messages are those appearing in either the CREATING or SETTING sections. Otherwise, the messages valid as create-time messages are defined by the type without any specific syntactic marker.

The createEnd message completes the process of specifying available options for an object being created. Typically it validates that requested options are valid and consistent with one another, and raises an error if they are not. The standard, predefined error InvalidCombination
(Error.html#stderr)
may be raised by createEnd to indicate an invalid combination of requests, or other, more specific forms of error handling may be used.

If all requests received since the initial createBegin: are valid, both individually and in combination with each other, then createEnd determines a finalized form of object that satisfies all requests received and then returns this object. Any additional storage required for the finalized object is taken from the same zone originally passed to createBegin. The object may have whatever implementation is selected to best satisfy a particular request. Different requests may result in entirely different implementations being returned. The only guarantee is that a returned object supports the messages defined for further use of the finalized object. If a type was defined by a @deftype declaration, these messages are those appearing in either the SETTING or USING sections.

On return from createEnd, the id of the interim object returned by createBegin: is no longer guaranteed to be valid for further use, and should no longer be referenced. A variable which holds this such an id can be reassigned the new id returned by createEnd, so that the same variable holds successive versions of the object being created. Following is a fragment that illustrates typical coding practice using createBegin: and createEnd:

```
newArray = [Array createBegin: aZone];
  [newArray setInitialValue: aList];
  [newArray setDefaultMember: UnsetMember];
  [newArray setCount: [aList getCount] * 2 );
  newArray = [newArray createEnd];   // ! note reassignment of newArray
```

By resetting the variable to the returned value of createEnd, there is no possibility that the id value first returned by createBegin: can ever be used again. Sometimes, the id value returned by createEnd might be the same value first returned by createBegin:, if createBegin: was able to allocate the same object returned later by createEnd. In other cases, however, the object returned by createBegin: might be strictly a temporary object that is thrown away after createEnd determines a final object to be created and returned. To protect against changes in implementation, only the final createEnd return value should ever be used again after createEnd has returned.

```
+ customizeBegin: aZone;
- customizeEnd;
- customizeCopy: aZone;
```

Some types accept create-time messages not only when creating a new instance, but to customize a new version of the type itself. Objects created from a customized type will have all options preset that create-time messages sent to the customized type object have already set. If many objects all need the same create-time options, it is often simpler (and can also be faster) to create a customized version of a type first, and then create further instances from that type.

Customizing a type object does not modify the original type object, but instead creates a new type object that has the customizations built-in. A create: message on the new type object creates a new instance as if the same sequence of create-time messages had been sent to the original type object using createBegin: and createEnd. A type is customized by bracketing the sequence of create-time messages not with the createBegin: and createEnd messages used to create a new instance, but with customizeBegin: and customizeEnd messages instead. customizeBegin: returns an interim value for receiving create-time messages much like createBegin:, and customizeEnd returns the new, customized version of the original type. Following is sample code sequence that illustrates the use of these messages:

```
newArrayType = [Array customizeBegin: aZone];
  [newArrayType setCount: 100];
  newArrayType = [newArrayType customizeEnd];
  array1 = [newArrayType create: aZone];
  array2 = [newArrayType create: aZone];
  ... // [array1 getCount] and [array2 getCount] are both 100
```

The customizeCopy: message creates a new copy of the interim object returned by customizeBegin: which may be used for further customizations that do not affect the customization already in progress. It may be used to branch off a path of a customization in progress to create an alternate final customization. Following is a code fragrment which illustrates this usage:

```
newArrayType1 = [Array customizeBegin: aZone];
  [newArrayType1 setCount: 100];
  newArrayType2 = [newArrayType2 customizeCopy: aZone];
  [newArrayType2 setDefaultMember: UnsetMember];

  newArrayType1 = [newArrayType1 customizeEnd];
  newArrayType2 = [newArrayType2 customizeEnd];
  array1 = [newArrayType1 create: aZone];  // no DefaultMember option
  array2 = [newArrayType create: aZone];   // DefaultMember option set
```

customizeCopy may be used only on an interim object returned by customizeBegin: and not yet finalized by customizeEnd. The new version of the interim object being customized may be allocated in the same or different zone as the original version, using the zone argument required by customizeCopy:

The zone passed to customizeBegin: is the same zone from which storage for the new, finalized type object will be taken. This zone need not be the same as any instance later created from that type, since a new zone argument is still passed in any subsequent create message on that type.

Whether a customized version of a type can be created depends on the implementation of the type itself. If a type does not support customization, a customizeBegin: message on the type raises an error. All types defined by an @deftype declaration may be relied on to support at least one cycle of customization to create a new type object. Whether an already customized type object (returned by customizeEnd) supports a further cycle of customization (by another sequence of customizeBegin:/customizeEnd) depends on the implementation of the original starting type. A type should not be relied on to support more than one cycle of customization unless it is specifically documented to do so.

Any interim object returned by either createBegin: or customizeBegin: supports the getZone and drop messages that a finalized instance may also support. These messages are defined by the Drop
(Drop.html)
type, which is normally inherited by a type to declare these messages on a finalized instance. This type is not inherited by the Create type because the messages would then apply to the finalized instance, not to the interim object. Even though not declared, the messages are available on the interim objects nonetheless. The drop message on an interim object may be used if it turns out that a finalized version is no longer required after creation or customization has already begun.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Drop – end existence of an object

Synopsis

```
//
// Drop -- end existence of an object
//
@deftype Drop
- (void) drop;
@end
```

Description

The Drop supertype defines the drop message, which is a standard message for indicating that an object no longer exists and will never again be referenced. Any future attempt to reference a dropped object is an error. This error may or not produce predictable effects depending on the level of debug checking and other factors.

Immediate effects of the drop message depends on the subtype of Zone (Zone.html) used to provide storage for the object. For some zone types, the drop message immediately deallocates storage for the object and makes the freed storage available for other use. Subsequent use could include the allocation of a new object at precisely the same location, resulting in a new object id identical to a previously dropped one. Other zone types (..none implemented yet..) might delay any reclamation of freed storage until a more comprehensive check can be made that other references no longer exist to the dropped object. When running in debug mode (see debug support (globaldef.html#debug)) the storage for a dropped object is filled with predictable contents so that an error is raised if there is any immediate attempt to us it again.

The Drop type may be inherited by any type that provides drop support for its instances. In addition to freeing the storage and invalidating the object, a drop message may release other resources acquired or held within the object. Not every object which can be created can also be dropped, and some objects can be dropped which are not directly creatable. Some objects may be created as a side effect of other operations and still be droppable, and some objects may be created with links to other objects and not droppable on their own. A type independently inherits Create or Drop types, or both, to indicate its support of these standard interfaces to define the endpoints of an object lifecycle.

---

*Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)*

Last modified: 1998-01-08

Name

Zone – modular unit of storage allocation

Synopsis

```
//
// Zone -- modular unit of storage allocation
//
@deftype Zone <Create, Drop, GetOwner, CREATABLE>
CREATING
- (void) setReclaimPolicy: reclaimPolicy;
- (void) setStackedSubzones: (BOOL)stackedSubzones;
- (void) setPageSize: (int)pageSize;
USING
- getReclaimPolicy;
- (BOOL) getStackedSubzones;
- (int) getPageSize;

- allocIVars: aClass;
- copyIVars: anObject;
- (void) freeIVars: anObject;

- allocIVarsComponent: aClass;
- copyIVarsComponent: anObject;
- (void) freeIVarsComponent: anObject;

- getComponentZone;

- (void *) alloc: (size_t)size;
- (void) free: (void *)aBlock;

- (void *) allocBlock: (size_t)size;
- (void) freeBlock: (void *)aBlock blockSize: (size_t)size;

-                getPopulation;
- getSubzones;
- (void) mergeWithOwner;

- (BOOL) containsAlloc: (void *)alloc;
- getSubzone: (void *)alloc;

- (void) reclaimStorage;
- (void) releaseStorage;
```

```
- (void) xfprint;
- (void) xfprintid;
@end
```

Description

A zone is a source of storage for objects or other allocated data. Whenever a new object is created, a zone must be identified from which the storage for its instance variables, or other internal data, is obtained. A program may establish multiple zones to ensure that objects with similar lifetime or storage needs are allocated together, and in general to optimize allocation and reuse of storage.

Zones also maintain a collection of all objects allocated within the zone. This collection, referred to as the "population" of a zone, is a set of all objects which have been created but not yet dropped within the zone. Collections maintained automatically by zones can eliminate a need for other, separately maintained collections in applications that need to keep track of entire populations of objects. Collections of allocated objects can provide support for object query, external object storage, and automatic storage reclamation.

A zone may be used to obtain storage not only for objects, but also for raw storage blocks like those provided by the C malloc function. All objects and storage blocks allocated in a zone remain local to that zone. This means that allocation of storage in other zones does not affect the efficiency of storage allocation within a particular zone. For most zone types, individual allocations may still be freed within a zone, and total storage of a zone may grow and shrink according to aggregate needs. In addition to freeing individual allocations, an entire zone may also dropped. Dropping a zone automatically frees all allocations made within it, including final drop processing on any allocated objects that need it. Release of an entire zone can be much faster than individual release of each object within it.

The Zone type is a fully implemented type that provides default storage management support for objects and other allocated storage. It is also a supertype for other zones that implement alternative policies for use in specialized situations.

(..The current Zone implementation is a temporary one based on a simple pass-through of all allocation requests to the C library malloc function. Complete support of the zone population is already provided, however, as is support for internal storage mapping that will eventually be required for automatic storage reclamation Considerable improvements in storage allocation are possible when the current implementation is replaced by more efficient allocation contained in pages local to each zone. ..)

Inherited behavior

```
@deftype Zone <Create, Drop, CREATABLE>
```

A zone is created using standard create messages just like other objects. This means that a zone must identify another zone from which it obtains its storage. Storage is typically obtained from this other zone in large units called pages, which are then managed by the local zone to support internal allocations. The getZone message of the Drop (Drop.html) type returns the zone which provides these base pages.

Since a new zone always requires that an existing zone be identified, no new zones could be created unless there were some zones that already existed. Several such zones are predefined as part of the defobj library; see predefined zones (globaldef.html#zones) for a summary.

Create-time behavior

```
- (void) setReclaimPolicy: reclaimPolicy;
- (void) setStackedSubzones: (BOOL)stackedSubzones;
- (void) setPageSize: (int)pageSize;
USING
- getReclaimPolicy;
- (BOOL) getStackedSubzones;
- (int) getPageSize;
```

ReclaimPolicy

(.. Not implemented yet. Will define the various types of zone implementations, all the way to zones that support full automatic storage reclamation. ..)

StackedSubzones

(.. Not implemented yet. Specifies that no further allocations are made in the zone after a subzone is created, until that subzone itself is dropped. Supports an especially efficient of implementation. ..)

PageSize

PageSize specifies the size of pages within which a zone manages its internal allocation. Its default is typically a natural page size (perhaps 4K) for the local machine architecture. The default should be overridden only when tuning storage allocation for very specific situations. Allocations within a zone are not limited to the page size, since any requests that exceed the page size are simply passed up to the owner zone within which the zone was allocated.

Object allocation

```
- allocIVars: aClass;
- copyIVars: anObject;
- (void) freeIVars: anObject;
```

allocIVars: allocates the instance variable structure for a new object. The initial word of this structure is set to class id passed as its argument. The class also determines the size of the structure allocated. All remaining contents of this structure are initialized to binary zeroes.

copyIVars: creates copies an existing instance variable structure into a new allocation made within the local zone. The existing instance variable structure may be in any zone, but must contain a class pointer in its first word that correctly describes the size of the structure.

freeIVars: releases storage that was previously allocated to hold the instance variable structure of an object. The first word of the object must be a class pointer that correctly describes the size of the structure. Storage allocated by allocIVars: or copyIVars: may be freed only by freeIVars:, and freeIVars: may be used only to free storage allocated by one of these messages.

The object allocation functions are intended for use only within the implementation of an object class. Users of an object type should create a new object only using create: messages or other messages provided for this purpose.

Internal component objects

```
- allocIVarsComponent: aClass;
- copyIVarsComponent: anObject;
- (void) freeIVarsComponent: anObject;

- getComponentZone;
```

These messages allocate, copy, and free the storage that holds the instance variables for an object, but they allocate the object as an internal component of the zone that is not included in the zone population. They are used by classes that allocate additional objects as part of the implementation of another object, and that control the mapping of this storage separately from the zone level objects. getComponentZone returns a specially qualified version of the zone that automatically allocates all its objects with the internal component qualification, even if allocated with allocIVars: or copyIVars:. This qualified zone may be passed as an argument to a create: or createBegin: message so that it will create the new object as an internal component object.

Raw storage block allocation

```
- (void *) alloc: (size_t)size;
- (void) free: (void *)aBlock;

- (void *) allocBlock: (size_t)size;
- (void *) copyBlock: (void *)aBlock blockSize: (size_t)size;
```

```
- (void) freeBlock: (void *)aBlock blockSize: (size_t)size;
```

alloc: allocates a new storage block much like the malloc function of the C library. The storage is aligned according to the most restrictive requirements for any data type on the local machine architecture. The storage is not initialized to any known contents.

free: releases a block of storage previously allocated using alloc:. The size of the block is not required as an argument because alloc: has saved this size as necessary as part of the initial allocation. free: may be used only to free a block allocated by alloc:, and a block allocated by alloc: may be freed only by free:.

allocBlock: allocates a new storage block similar to alloc:, except that the size of the block allocated must be passed as an argument when freeing the block. copyBlock:blockSize: allocates a new storage with the same requirements on free as allocBlock:, but also fills the new storage with the contents of a block passed as the first argument, up to the size of the newly allocated block. freeBlock:blockSize: must be used to free any block previously allocated by allocBlock: or copyBlock:blockSize.

Subzone structure

```
-               getPopulation;
- getSubzones;
- (void) mergeWithOwner;

- (BOOL) containsAlloc: (void *)alloc;
- getSubzone: (void *)alloc;
```

getPopulation returns a collection all objects allocated in a zone using either allocIVars: or copyIVars: and not yet freed using freeIVars:. getObjects returns nil if the ObjectCollection option is false. The collection returned has the type OrderedSet
(../collections/OrderedSet.html)
as defined in the collections library
(../collections.html)
, with the ReadOnly option set true and the IndexSafety option set to SafeAlways. The members of this collection may change as objects are allocated and freed, but may not added or removed directly within the collection.

getSubzones returns a collection of other zones that all obtain their storage from the local zone. mergeWithOwner drops the local zone, but reassigns all the zone allocations to the zone from which all local storage was obtained. (..neither message currently implemented..)

containsAlloc: tests if a particular allocation was made by the local zone, or any zone which obtains storage from that zone. The pointer argument must point to an allocation previously made by any zone. The message returns true if the allocation was made by either the local zone or any of its subzones. getSubzone is similar, except that it returns the particular zone in which the pointer was allocated, or nil if the pointer was not allocated in the local zone or one of its subzones. (..neither message currently implemented..)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Symbol – object defined as a distinct global id constant

Synopsis

```
//
// Symbol -- object defined as a distinct global id constant
//
@deftype Symbol <Create, GetName, CREATABLE>
CREATING
+ create: aZone setName: (const char *)name;
- (void) setName: (const char *)name;
@end
```

Description

A Symbol is an object created with a fixed name. It has no behavior except to get the name with which it was created. A Symbol is typically used to define unique id values which are assigned to global constant names. These names, capitalized according to the recommended convention for global object constants, are used by some libraries as flags or enumerated value codes in arguments or return values of messages.

Ordinarily, a symbol is created with its character string name matching the global id constant to which it is assigned. These global program constants can then provide a minimal level of self documentation as objects. Subtypes of Symbol can extend the base of a named, global id constant to establish further components of a global, constant definition.

Inherited behavior

A symbol is fully creatable using standard Create messages. A character string name must be supplied for any new symbol; there is no default. Symbol inherits the getName

(misc.html#GetName)
message, which returns the symbol name.

Create combination messages

```
CREATING
+ create: aZone setName: (const char *)name;
```

create:setName: is a combination message defined as a caller convenience. See combination messages
(combo.html)
for a summary of conventions on combination messages.

Create-time options

```
CREATING
- (void) setName: (const char *)name;
```

The Name option may be set only at create time. Its value is a null-terminated character string that remains fixed for the life of the object. The inherited getName message returns this name from a created instance.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Warning, Error – report of some condition occurring during program execution

```
(..  This section of documentation is only a start.)
```

Synopsis

```
//
// EventType -- report of some condition occurring during program execution
//
@deftype EventType <Symbol, RecordType>
- (void) raiseEvent;
- (void) raiseEvent: (void *)eventData, ...;
@end

//
// Warning -- a condition of possible concern to a program developer
//
@deftype Warning <EventType, CREATABLE>
- (void) setMessageString: (const char *)messageString;
- (const char *)getMessageString;
@end

//
// Error -- a condition which prevents further execution
//
@deftype Error <Warning, CREATABLE>
@end
```

Description

```
Warning and Event are object types (both subtypes of EventType) that may be used to report
conditions detected during program execution.  A warning is a condition of possible interest
to a developer or user, but one that does not prevent further execution.  An error is a condition
which prevents further execution of the program.
```

```
Warning and Event also inherit from Symbol, so that each type of warning or error
may be given a unique global name.  Any library may publish its own warning and event
names as part of its documented library interface (see module definition conventions
```

(module.html)
for standard methods of declaring interface symbols in a library header file.)

These predefined warning and event types are typically used not directly by the messages defined above, but by the following macro defined in defobj.h:

```
#define raiseEvent( eventType, formatString, args... )
```

This macro raises an event of the type identified by the first argument. Default handling of the event is to print a formatted message including standard header and trailer lines, and then either to continue (in the case of a Warning) or to terminate the program in the case of Error). The standard header lines include the file, function, and line number where the macro that raised the event was executed.

If the formatString argument is nil, a default message string already established as part of the warning or event type is used to report the error. Otherwise, the formatString argument must be a null-terminated C character string that overrides the default string.

Either the default or replacement message string may contain printf-style formatted field definitions. For each formatted field definition, an argument value must be supplied in the args... section of the macro. (The support of the GNU C compiler for macros with variable argument lists is used by the macro.)

(.. Methods for defining new event types as part of a library implementation are not currently documented. Instead, just use the raiseEvent macro on one of the predefined Warning or Event Types if you want to raise an event using this standard mechanism. Eventually, this standard mechanism will be used to handle warnings and errors under various forms of logging exception handling support.)

Standard Error and Warning Types

The following standard error and warning types are defined by the defobj library. Thesy may also be used by other libraries for errors or warnings that fit within these standard predefined categories:

```
//
// standard errors
//
extern id <Error>
  SourceMessage,          // message in the source defines error
  NotImplemented,         // requested behavior not implemented by object
  SubclassMustImplement,  // requested behavior must be imple-
mented by subclass
  InvalidCombination,     // invalid combination of set messages for create
  InvalidOperation,       // invalid operation for current state of receiver
  InvalidArgument,        // argument value not valid
  CreateSubclassing,      // improper use of Create subclassing framework
  CreateUsage,            // incorrect sequence of Create protocol messages
```

```
  OutOfMemory,            // no more memory available for allocation
  InvalidAllocSize,       // no more memory available for allocation
  InternalError,          // unexpected condition encountered in program
  BlockedObjectAlloc,     // method from Object with invalid allocation
  BlockedObjectUsage;     // method inherited from Object superclass

//
// standard warnings
//
extern id <Warning>
  WarningMessage,         // message in the source defines warning
  LibraryUsage,           // invalid usage of library interface
  DefaultAssumed,         // non-silent use of default
  ObsoleteMessage,        // using feature which could be removed in future
  ObsoleteFeature,        // using feature which could be removed in future
  ResourceAvailability;   // resource from runtime environment not available
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1997-01-08

Miscellaneous Operations

Object Types

The following object types are documented below:

- Copy
- GetName
- GetOwner
- SetInitialValue

Summary of Miscellaneous Operations

An assortment of object types in the defobj library serve just to establish a single declaration of messages that many other object types adopt as part of their declared interface. The centralized definition of the messages in these types avoids their repetition throughout many other types. It also provides a single place where basic rules for these messages can be summarized for all other implementations to follow. This reference documentation provides these summaries along with intended purpose of each message.

> None of these types inherits from any other types; their whole purpose is just to package the definition of one or more messages. The type name is typically the same as a message declared by the type, only with the uppercase convention of a type name. These types may be freely mixed into the inheritance lists of other types to build whatever combinations of messages those other types need to define.

Copy Interface Description

Synopsis

```
//
// Copy -- copy all state defined as part of object
//
@deftype Copy
- copy: aZone;
```

```
@end
```

Description

The copy message creates a new object that has the same contents and resulting behavior as a starting object, except that an independent copy of the contents of the starting object is created so that further changes to one object do not affect the other. The zone argument specifies the source of storage for the new object. The message returns the id of the new object created.

> An object type that supplies the copy operation defines what it includes as the contents of an object copied. There is no global rule for what is considered "inside" a copied object vs. merely referenced by it. (There is no fixed notion of "shallow" vs. "deep" copy found in some object libraries.) After copying, the new object may still contain some references to other elements also referenced by the starting object, but in general the new object minimizes any dependencies shared with the starting object. Any object type supplying the copy message should also supply documentation on its rules for copied objects.

GetName Interface Description

Synopsis

```
//
// GetName -- get name which identifies object in its context of use
//
@deftype GetName
- (const char *) getName;
@end
```

Description

The getName message returns a null-terminated character string that identifies an object in some context of use. This message is commonly used for objects that are created once in some fixed context where they are also assigned a unique name. Constant objects defined as part of a program or library are examples. This message is intended only for returning a name associated with an object throughout its lifetime. It does not return any data that ever changes.

GetOwner Interface Description

Synopsis

```
//
// GetOwner - get object on which existence of object depends
//
@deftype GetOwner
- getOwner;
@end
```

Description

The getOwner message returns another object which is considered as the owner of an initial object. What is considered as an owner depends on its specific object type, but might be a larger object of which the local object is a part, or an object that has exclusive control over the local object. The principal constraint established by an ownership structure is that a given object can have only a single other object as its unambiguous owner.

Ownership hierarchies arrange themselves in a strict, single-rooted tree. The top-level node of an ownership hierarchy typically returns nil as its owner. If an object is regarded merely as one part of another object defined as its owner, then copying or dropping the owner object should copy or drop the member object as well. Owner and member are neutral terms for a generic relationship sometimes called parent vs. child, but it is up to a particular object type to define specifically what it means by a getOwner relationship.

SetInitialValue Interface Description

Synopsis

```
//
// SetInitialValue - create using initial value from an existing object
//
@deftype SetInitialValue
CREATING
+ create: aZone setInitialValue: initialValue;
+ create: aZone setReadOnlyValue: readOnlyValue;

- (void) setInitialValue: initialValue;
- (void) setReadOnlyValue: readOnlyValue;
SETTING
- (void) setReadOnly: (BOOL)readOnly;
USING
- (BOOL) getReadOnly;
@end
```

Description

The SetInitialValue type defines a variety of messages relating to an initial or unmodifiable value established as part of an object. This message is typically provided when creation of a new object might be more easily accomplished by copying the value of an existing object rather than establishing a new value from scratch. As with the copy message, precisely what is considered the value of an existing object to copy is defined only by the particular object type that supplies these messages.

If an object has a value which can be established at create time, it is often useful (and can also enable significant optimization) to declare that no further modification will occur to this value during further use of the object. A restriction against modifying a value is referred to as a "read-only" restriction. This type supplies messages to declare a read-only restriction along with any initial value. For some object types, a read-only restriction can also be added or removed after an object has already been created.

The setInitialValue: message requires another object as its argument, from which the value of a newly created object is to be taken. Unlike a copy message, the object used as the source of the new value need not have the identical type as the new object to be created. A particular object type defines the types of initial value objects which it can accept, along with any special conversion or interpretation it might apply to such a value.

The ReadOnly flag specifies that the value of an object may not be modified further by other messages on an object. If such messages are attempted then an error is raised. If the ReadOnly flag is set at create time, the object is first created in read-only mode. Typically an initial value is always provided for an initial read-only object, since otherwise the created object would have only an empty or initial default value. The setReadOnlyValue: message is equivalent to a combination of a setInitialValue: message with an added ReadOnly restriction.

If a read-only restriction is established at create time, then for some object types this restriction may be required to remain for the entire lifetime of an object. For other object types, the read-only restriction may be added or removed any time after initial object creation. If a particular type does not support a change to read-only status, a request to do so raises an error.

Standard forms of create combination messages
(combo.html)
are provided to create an object in one message expression from the value of another object, with or without a read-only restriction.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Module Definition Conventions

Overview

Some of the basic Swarm libraries are implemented using object definition conventions established by the defobj library

(../defobj.html)
. These conventions include some minor special syntax to declare the public interfaces to Objective C objects. This document explains how to read and interpret the header files and interface definitions published by such libraries. It also explains the typical structure of documentation provided for libraries that follow these conventions.

These conventions are documented under the following subsections:

- library header file
- object type definitions
- GridTurtle example
- global object symbols
- interface design conventions
- documentation structure

Library Header File

For a library that adopts these conventions, a library is not merely a collection of source files that are compiled into a library archive under control of a make file. Instead, Objective C source files in the library are processed in a special way to publish their definitions not only as header files, but as generated objects available at runtime to make full use of the library. A library processed in this way is referred to as as a "module."

The defobj
(../defobj.html)
library documents the full details of special processing performed on a module. For simply using a library defined as a module, the key fact to keep in mind is that the entire public interface to a module is declared in the one header file having the same name as the module itself, plus a trailing .h suffix. Additionally, the header file of a library module normally documents only the public interface to a module, in a way that is completely separated from the implementation of the objects it specifies.

The separation of implementation means that a library module publishes its interface entirely without reference to any Objective C classes which implement its objects.

Even though classes are often thought of as separating the interface of an object from its implementation, this separation is far from complete. Not only do classes typically contain many internal methods not intended for external use, but they also define a particular storage format for an object defined as instance variables.

A library module instead publishes its interface as a set of public object types. These object types may also be supplemented by global object constants called symbols. Both these kinds of definitions normally appear only in the header file of a library. The remaining source files in a library normally contain the classes which implement the object types.

In a library module, the files which implement classes (including class header files), need not ever be referenced simply to make use of the implemented capabilities of the library object types. Documentation for the library is normally expressed entirely in terms of the types and symbols published in the library header file. If a feature does not appear in the library header file, it should not be considered part of a supported public interface.

Individual class header files are required to subclass from existing implementations, but interfaces for subclassing are an entirely separate issue from normal public use of an object library. Class inheritance can be a powerful implementation technique, but extension of an existing class framework is typically safe only if performed in explicitly permitted ways. If a library supports subclassing at all, it must carefully state which classes may be subclassed and in what ways. For a library module, this information is supplied outside the library header file. The library header file specifies only the interfaces by which objects are intended to be used, whether implemented by a local class or an external subclass.

Remaining sections of this document explain the declarations which appear in a library header file, and end with a suggested structure of documentation to be provided for a library module. The library modules of Swarm mostly follow this structure.

Object Type Definitions

The interfaces to objects defined by a library module are specified by object type definitions. An object type defines only a set of messages which may be sent to an object over various phases of its lifetime. An object type makes no commitment to the classes that might be used to implement the object.

Multiple classes may all implement the same messages belonging to a type. The independence of types and classes means that different classes can provide alternate implementations of the same object type. For example, a particular implementing class might be selected to optimize the implementation for a particular case.

Object types are similar to protocols defined by Objective C, and the declarations appearing in a library header file are a minor adaptation of Objective C protocol syntax. A key difference from protocols is that object types are published as real external objects that may be used at runtime to create instances of a type. A futher difference is that the object types of the defobj library are divided into separate interfaces that define distinct phases of an object's life cycle.

GridTurtle example

An example will help illustrate the features of type definitions supported by the defobj library. Throughout documentation of the basic Swarm libraries, a series of running examples will be based on a simple type of object belonging to sample Swarm simulations. This type of object is an agent that move around on a two-dimensional grid, always moving in a current direction that it maintains internally. This agent is like a "turtle" of the original Logo system, except that its position is constrained to discrete integer values of its X-Y coordinates, and its direction is always one of the four orthogonal directions north, east, south, or west.

Following is a complete library header file for a module which defines such an object, called GridTurtle:

```
/*
Name:         GridTurtle.h
Description:  object type for Swarm example programs
Library:      grid
*/

#import <defobj.h>
@deftype GridTurtle <Create, Drop, CREATABLE>CREATING
- (void) setXLocMax: (int)xLocMax;
- (void) setYLocMax: (int)yLocMax;
SETTING
- (void) setXLoc: (int)xLoc;
- (void) setYLoc: (int)yLoc;
- (void) setXLoc: xLoc setYLoc: yLoc;

- (void) setDirection: direction;
USING
- (int) getXLoc;
- (int) getYLoc;

- getDirection;

- (void) move: (int)distance;
- (void) turn: (int)angle;  // angle measured in units of pi/2 (90 deg.)
- (void) print;
@end

id <Symbol> North, East, South, West;
```

```
#import "grid.xt"
```

An object type is defined by an @deftype declaration. The syntax of such declaration is identical to that of an Objective C @protocol definition, except for the entirely uppercase keywords (CREATABLE, CREATING, SETTING, USING) appearing in the GridTurtle example above. All these modifications of Objective C syntax are accomplished by simple preprocessor macros; no extensions to the language compiler are involved.

When this library header file is processed (by a special rule in a make file), an external object id with the name GridTurtle is automatically published. The name of a defined type becomes an ordinary object that accepts specific messages defined by the defobj library. The defobj library
(../defobj.html)
explains the details of such messages; the only purpose here is to explain the basic sections of a deftype declaration.

deftype declarations follow the syntax as Objective C protocols for inheriting messages from each other: a list of names enclosed in angle brackets (e.g., <Create, Drop, ...> above) gives the names of other declared types containing messages to be supported by the new type as well. (These types referenced here are defined by the imported file <defobj.h>.) Like protocols, full multiple inheritance of types is supported. The same messages may be inherited any number of times through any path with no different effect than if inherited or declared just once, so long as no conflicts occur in any of their argument or return types.

The CREATABLE tag appearing in the inherited type list above is a special type which defines no messages of its own, but merely marks the type as one which supports direct creation of instances of the type. Without this tag, the only role of a type is to define messages for inheritance by other types. With this tag, the global type object has a particular implementation that supports object creation using standard messages defined in defobj.

The declared messages of the type may be separated into sections marked by the special uppercase tags such as CREATING, SETTING, and USING above. (Currently, these are the only such tags which may occur.) These sections each define messages belonging to a particular defined "interface" of the object type, which are further combined into distinct "phases" of an object lifecycle supported by defobj messages. Further explanation of the interfaces and phases defined by this example are provided in the Usage Guide
(guide.html)
of the defobj library.

The following line at the end of the <grid.h> header file above is required as part of the standard conventions for coding a module header file: This line adds no information to the header file, but is required as part of the system by which definitions are automatically generated.

```
#import "grid.xt"
```

Global Object Symbols

The <grid.h> header file above also contains the declaration:

```
id <Symbol>  North, East, South, West;
```

Lines that declare global id variables of type Symbol, EventType, Warning, or Error (using the angle bracket syntax of id variables conforming to a protocol) are processed somewhat like deftype declarations in that they also produce global id variables initialized to support particular messages defined by defobj. These global variables, however, are not used to define or implement other message interfaces, but only to define certain fixed capabilities referenced through their global object names.

If declared as a Symbol, as in the case here, the generated objects have no particular behavior of their own (other than the character string of their name), but only serve to define unique global id constants which may be used as distinct named values in messages. In this example, the current direction of a GridTurtle object is represented by one of the symbol names North, East, South, or West. These values are like the enum constants of the C language, except that they are defined as full Objective C objects, and may be used with further restrictions. An EventType, or a Warning or Error, defines a further subtype of a Symbol constant with further specialized messages documented in defobj.

Interface Design Conventions

A variety of rules on naming and declaration of object types, symbols, and messages are followed by many of the Swarm libraries. These rules help establish a basic consistency on the library interfaces. Following is a list of such conventions that apply to a public library interface, some but not all of which are derived from standard Smalltalk or Objective C coding practice:

- Names of global object constants are capitalized. In the public interface, such names include types and symbols.
- Recapitalization separates words of a compound name (e.g., GridTurtle). Underscores are generally not used.
- Message names start with a lower-case character, and are named using verbs. Nouns that represent gettable or settable components of object state (e.g., Direction of a GridTurtle), are prefixed by get or set to indicate the action being performed.
- The Smalltalk convention that a message return the receiver of a message if there is no other specific return value is generally *not* followed. If there is no specific return value needed from a message the return type is declared (void).

Documentation Structure

The following standard sections of documentation are suggested for libraries that follow the strict interface vs. implementation separation of library modules, each with purpose and typical contents as given:

- Interface Reference
  Complete, concise summary of all features of a library. Fills the role a Unix "man page" in providing comprehensive definition of library services. Little or no concern to provide path of tutorial introduction. Does not contain extended examples, at most only fragments of examples which serve needs of specification.
- Usage Guide
  Tutorial introduction to a library. Focuses on the most common uses in an order reflecting the needs of a first-time user. Most explanation by means of progressively elaborated examples. Serves as a guided tour of major library capabilities. No attempt at reference-style completeness.
- Advanced Usage Notes
  Continues overview of all significant capabilities of a library, including those which might be needed by advanced users customizing or extending built-in capability. Provides examples of specialized uses extending beyond normal, basic usage, but still using built-in features of the library framework. Helps to simplify the Usage Guide by providing a location for overflow of more advanced features.
- Subclassing Notes
  Documents the rules for writing new classes which subclass from classes that implement the library. Because a library might use complex combinations of classes to implement the range of behaviors defined by its types, the classes in a library are not automatically usable as superclasses of user-defined classes. Each library documents which classes are available for use as superclasses, and the specific rules that must be followed when subclassing from these classes.
- Interface Design Notes
  Explains why choices were made the way they were in the design of the public interface to a library, including other alternatives considered. Serves as a "Rationale" document for the documented interface. May include historical background deliberately left out of the preceding sections. May also include references to other work.
- Implementation Notes
  Explains the structure of classes by which the types of a library are implemented. Summarizes the status of implementation if still incomplete and lists items of possible future work. Provides overview and high-level structure of the implementation in whatever ways would best guide a reader of implementation source code. May also discuss tradeoffs considered along with references to related or supporting work.

Marcus G. Daniels <mgd@santafe.edu>

(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Create Combination Messages

The create phase of an object type may contain messages that are defined simply by identifying them as a combination of other defined messages. Such messages are referred to as "create combination messages." Typically, such messages have the following form:

- create: aZone setOption: optionValue;

For this message to be a valid create combination message, the object type must already have the create: and setOption: messages defined (either locally or by inheritance). The meaning of this message is then identical to the following sequence (where SomeType is the name of a type implementing the combination message):

- anObject = [SomeType createBegin: aZone];
- [anObject setOption: optionValue];
- anObject = [anObject createEnd];

Create combination messages are provided as a convenience for commonly used options. Some combination messages also have their own custom implementing methods so that creation is faster when using the combination message. Any such implementation, however, is required to implement the parts of the combination message so that its meaning is identical to the results that would be obtained if the separate parts were requested in separate message.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Global Portability Assumptions

In a pure object-oriented system such as Smalltalk, the type of an object reference (an object id in Objective C) is the only fundamental data type used within the system. In a hybrid object system such as Objective C, both object id's and other data types are used together. The Swarm libraries make some basic assumptions about the compatibility of native C data types with the id pointer type of Objective C. These assumptions allow some of the more generic abstractions of the libraries (such as collections and actions) to provide one fundamental abstraction that can still be used under an identical implementation to support most of the other C data types.

The basic assumption is that all data types except double may be cast to and from a stored id type without losing the original value. This assumption is valid on all modern 32-bit or greater machines, including all machines on which Swarm is currently running or planning to run.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Global Definitions

Contents

The following global definitions are documented below:

M() macro

Synopsis

```
//
// M() -- macro to abbreviate @selector()
//
#define M( messageName )  @selector( messageName )
```

The M() macro is currently nothing more than an abbreviation for the Objective C syntax used to obtain a selector for a particular message name. It is provided as a convenience for use in messages that combine a selector with other arguments.

func˙t type definition

Synopsis

```
//
// func_t -- function pointer type
//
typedef void (*func_t)( void );
```

func_t provides a type definition for a generic function pointer type (taking no arguments and returning no value). This type name is used in declaring messages that accept or return function pointers of various alternative types.

---

initModule() function

Synopsis

```
//
// initModule() -- module initialization macro
//
#define initModule( module )
```

The initModule function must be called to initialize all the definitions of a library module before any attempt is made to use the library. Initializing a particular library module will also initialize all other library modules on which the library depends. The macro has no effect if the library has already been initialized, so that multiple libraries that all depend on another library can all request initialization of the other library with no ill effect. The argument of the macro is the name of module to be initialized. For example, if just the collections library were being initialized, the macro would be coded:

Global Zones

Synopsis

```
#define globalZone
#define scratchZone
```

The identifiers globalZone and scratchZone refer to two standard zones created by initialization of the defobj library. All other created zones must obtain their storage from one or the other of these two initial zones. (These initial zones are defined as macros only to follow conventions on global external names.)

> globalZone may be used for all storage which must be allocated for the duration of the program, or until explicitly released. scratchZone should be used for storage which is allocated only for use during execution of a single function. It is typically used for objects which are strictly temporary as part of a local computation. (.. Currently, a scratch zone has the identical implementation as the global zone, but eventually a scratch zone may be optimized for temporary allocations, and such allocations may also be freed automatically at various points when functions are guaranteed to be complete.)

Output Files

Synopsis

```
extern FILE *_obj_xerror;     // output file for error messages
extern FILE *_obj_xdebug;     // output file for debugging messages
```

The file pointers _obj_xerror and _obj_xerror define standard output files used for error messages and debugging messages, respectively. If not set otherwise by the time of defobj module initialization, they both default to the standard error output defined by the C library.

Debug Support

Synopsis

```
extern BOOL _obj_debug;       // if true then perform all debug error check-
ing

extern void xsetname (id anObject, const char *name);  // debug set dis-
play name

extern void xprint (id anObject);              // debug object print
extern void xprintid (id anObject);            // debug object id print

extern void xfprint (id anObject);             // debug foreach object print
extern void xfprintid (id anObject);           // debug foreach id print

extern void xexec (id anObject, const char *name);  // debug method exec
extern void xfexec (id anObject, const char *name); // debug fore-
ach method exec
```

The _obj_debug variable enables debug-level error checking in libraries that support optional levels of checking. Its default value is true. The default may be overridden either before initialization of the defobj module, or by means of a compiler macro setting (-D_obj_debug=0). If set to zero by a compiler macro setting, optimized compilation of code that refers to it will see only the constant zero and so avoid even generating code to test the variable.

The xsetname function sets the display name for an object using its setDisplayName: message. This string may be used to identify the object in a user-meaningful way; if defined it will appear in the identification string generated by a standard describe: message.

The xprint function prints a standard object description to the standard debugging output file. It generates the string to be printed using the describe: message defined by the DefinedObject supertype, but it is defined as a function so that it may be called directly from a debugger. xprintid is similar, but it prints only a one-line description using the built-in default.

The xfprint function is similar to the xprint function, except that it prints a debug description for each member of a collection. The argument of xfprint must support enumeration of collection members using the messages of the Collection
(../collections/Collection.html)
type. xfprintid is similar, but only one line per member is printed.

The xexec function executes a message on an object passed as an argument. The name of message to be executed (the full name including arguments as would be used in naming a selector) must be given as the second argument. This function is equivalent to the "perform:" message of Objective C, but is defined as a function so that it may be called directly from a debugger.

The xfexec function is similar to the xexec function, except that it executes a message on each member of a collection. The first argument of xfprint must support enumeration of collection members using the messages of the Collection
(../collections/Collection.html)
type.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Subclassing Reference

The defobj library currently includes the CreateDrop and Object_s superclasses used by other libraries. For user objects in Swarm simulations the SwarmObject superclass in the objectbase
(../objectbase.html)
library packages all required behavior of these superclasses into a single, simpler superclass.

The defobj library also provides support for custom-generated classes, such as those which implement separate phases of the standard create protocol. (See generated phase classes
(adv.html#phase-classes)
in the Advanced Usage Guide
(adv.html)
.) Subclassing from custom-generated classes is still not officially supported, because the framework for custom class generation is still being finalized. Custom-generated classes are currently used only in the defobj, collections, and activity libraries; these libraries document the specific classes that provide subclassing support.

In general, classes that implement a library are not automatically available for use as superclasses. Any library must document those specific superclasses that are valid for user classes to subclass from, along all the rules that a user subclass must follow when implementing new behavior.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Interface Design Notes

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Defobj Library: Implementation Notes

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library

Documentation and Implementation Status

Overview

The object types of the collections library establish a general-purpose foundation to maintain object references or other values as members of structured collections. Customization options consolidate a wide range of basic collection structures into a few core types (Array, List, Set, Map). These types are defined strictly by their interface, not their internal implementation. Specialized options, however, give control and flexibility for efficient, low-level use (such as implementing other libraries).

Documentation

The collections library follows the library definition conventions

(defobj/library.html)
. It has following standard sections of documentation:

- Usage Guide
(collections/guide.html)
- Advanced Usage Guide
(collections/adv.html)

- Interface Reference
(collections/ref.html)
- Subclassing Reference
(collections/subclass.html)

- Interface Design Notes
(collections/design.html)
- Implementation Notes
(collections/impl.html)

Marcus G. Daniels <mgd@santafe.edu>

(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Documentation and Implementation Status

The collections library has an almost complete set of Interface Reference documents, and the design of the interface is almost entirely final. The implementation of the collections library, however, doesn't yet implement many of the more advanced features of this interface.

Collections library development has concentrated on the specific features needed by the activity library. This library uses collections for its underlying support of actions to be executed in schedules. The collections library itself is still being completed for use as a general-purpose library.

The Array and List types have all basic capability fully implemented. Set and Map have basic messages defined, but currently rely on a crude implementation based on sorted lists. (This implementation is more than adequate for the usually small and relatively static schedule structures in the activity library, and is the most forgiving when change does take place within members being traversed.) Much more efficient implementations based on both balanced trees and hash tables are in the works.

OrderedSet is currently supported only with the low-level option for an internal member slot, and the implemented messages do not match the ones documented in the Interface Reference. Support for groups of duplicate members is missing from both Set and Map.

None of the special options for restricted usage modes on any type of collection (e.g., read-only restriction) has been implemented. Stack and Queue are not implemented, but are nothing more than restricted uses of a List. There is no support for any member type except id or smaller (other member types will depend on a data type facility to be supplied by defobj).

Even though incomplete, the portion of capability that is implemented has been exercised very heavily. The interfaces to Set and Map structures will remain the same even as their underlying implementations improve, so there is no harm in using them. See the GridTurtle test programs
(../grid.html)
for the most complete examples of collections usage.

The collections library follows the documentation structure suggested by the library interface conventions
(../defobj/library.html)
of the defobj library. There are placeholders at least for each section of documentation (some of which merely indicate that the section is not available yet) so that all links should at least link up with something, whether or not there's anything there.

Throughout the documentation, a parenthesized comment that starts with (.. indicates an editorial comment on the current status of implementation or documentation.

The documentation priority for all libraries is to complete at least their interface reference documents, so that there is the equivalent of Unix "man pages" that summarize all basic capability. A second priority is to complete the complementary "Usage Guide" documents. Unlike the reference documents, the Usage Guide will have a task-oriented organization, and will lead the initial user through actual code examples in the rough order a typical user is likely to need them. It will serve the role of a tutorial on each library.

The Usage Guide code examples have not yet been developed. For the time being, a directory of test programs (GridTurtle test programs
(../grid.html)
, contained within the documentation release directory) provides code examples of many of the basic features of the defobj, collections, and activity libraries. These code examples also help indicate the portions of the libraries which are fully implemented and working, since they are run on each new release of these libraries.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Usage Guide

This section of documentation is not yet available. In the meantime, see the GridTurtle test programs
(../grid.html)
for the most complete examples of collections library usage. If you can't find an example there that exercises a message or option that you want to use, chances are it's not implemented.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Advanced Usage Guide

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Interface Reference

Dependencies

Following are the other header files imported by <collections.h>:

```
#import <defobj.h>
```

The collections library follows the library interface conventions
(../defobj/library.html)
of the defobj library
(../defobj.html)
. It also depends on standard supertypes and classes defined by this library. Initialization of the collections library automatically initializes the defobj library as well. Since defobj also requires the collections library, both must always be linked into an application together.

Object Types (with links to reference pages)

```
Collection
 (Collection.html)
          String
 (String.html)
      OutputStream
 (Stream.html)
   Index
 (Index.html)

     /       |         \
    /        |          \
 Array
 (Array.html)
     List,
 (List.html)
           KeyedCollection,
 (KeyedCltn.html)

             ListIndex
 (List.html)
       KeyedIndex
 (KeyedCltn.html)
```

```
   /       \        /           \
Stack
 (Stack.html)
    Queue
 (Stack.html)
    Set,
 (Set.html)
         Map,
 (Map.html)

            SetIndex
 (Set.html)
     MapIndex
 (Map.html)

                                  |
          OrderedSet
 (OrderedSet.html)
```

(need a graphic with image map, but this will do for now)

Contents:

- Map
(Map.html)
- String
(String.html)
- OutputStream
(Stream.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Collection - supertype of all collection types

Synopsis

```
//
// Collection -- supertype of all collection types
//
@deftype Collection <Create, SetInitialValue, Copy, Drop>CREATING
- (void) setMemberType: aDataType;
- (void) setIndexSafety: indexSafety;
- (void) setIndexUpdateHandler: (func_t)handler : (void *)arg;
- (void) setReplaceOnly: (BOOL)replaceOnly;
USING
- getMemberType;
- getIndexSafety;
- (func_t) getIndexUpdateHandler: (void *)arg;
- (BOOL) getReplaceOnly;

- (int) getCount;

- atOffset: (int)offset;
- atOffset: (int)offset put: anObject;

- getFirst;
- getLast;

- (BOOL) contains: aMember;
- remove: aMember;
- (void) removeAll;

- begin: aZone;

- (void) forEach: (SEL)aSelector;
- (void) forEach: (SEL)aSelector : arg1;
- (void) forEach: (SEL)aSelector : arg1 : arg2;
- (void) forEach: (SEL)aSelector : arg1 : arg2 : arg3;
@end
```

Description

A collection is a grouping of object references or other data values which are assigned explicitly as members of the collection. Depending on the subtype, collection members may also be maintained in various associations with each other, such as an ordering of members or an association of members with identifying key values. Major Collection subtypes include Array, List, Set and Map. The Collection supertype establishes common conventions (and associated messages) supported by all types of collections.

All collections support traversal over their members using a separate object called an Index. All collections also follow common rules regarding the types of data values which may be added as members. The next two subsections summarize these basic features of all collections.

Index objects

An index is a special type of object that references a current position in an enumeration sequence over a collection. An enumeration sequence contains every member of a collection exactly once. Every collection defines an associated type of object for its index. The index type of a collection defines additional ways in which the members of a collection may be processed beyond any messages available directly on the collection. Often the operations of an index provide the most basic and flexible means for manipulating individual members of a collection. An index object into a collection may be created at any time. An index is the basic means to traverse through all members of a collection. Multiple indexes on the same collection may all exist at the same time and may reference the same or different positions. Depending on the collection type, it may be possible to modify a collection through its indexes.

Once an index is created, the sequence of members in its enumeration sequence is guaranteed to remain the same, provided that no new members are added to the underlying collection, or existing members removed. If a member is located once at a particular position, it is guaranteed to remain at that position as long as the index itself remains.

Many collection types define an explicit ordering over their members. For such collections, the sequence of members referred to by an index will always be consistent with this ordering. An explicit, total ordering also guarantees that all indexes of the same collection have the same member sequence.

If no such ordering is defined, however, some particular sequence of all the collection still becomes associated with each created index. All collection members are guaranteed to be contained somewhere in the enumeration sequence for any particular index, but two indexes on the same collection are not guaranteed to have the same sequence.

The Index type corresponds to the iterator types defined as part of many other object-oriented libraries. The name Index is shorter and emphasizes the more abstract and multi-function role of these basic support objects for any collection. For more background on design of indexes and iterators, see the Interface Design Notes for the collections library.

Member types

The default member type of any collection is the id type defined by Objective C. All messages which add, remove, or access members declare their member value with the id type. None of the collection types, however, requires that the member value actually point to a valid object. The member values of any collection may contain an arbritrary bit pattern, provided these bits fit in the space available within an id pointer.

In the few cases when default operations assume that a member points to a valid object, there are alternate ways to avoid these defaults and use an arbritrary bit pattern as a member value instead. When using an arbritrary bit pattern, the member value must be cast to or from the id type when adding or retrieving it from the collection. By global portability assumptions (../global-port.html) , the number of bits in an id value may be safely assumed to be at least as large as an int (signed or unsigned), but not necessarily as large as a long.

There is also an option to override the default member type so that it can be larger than a id pointer. Use of this option, however, changes the way member values must be passed to or from a collection. More information is given with the MemberType option below.

Inherited behavior

```
@deftype Collection <Create, SetInitialValue, Copy, Drop>
```

(.. SetInitialValue and Copy behavior is currently implemented for Array and List, but not for Set or Map.)

The Collection type inherits standard messages from the defobj library for creating and customizing objects. These messages are defined as part of the defobj Create supertype. The only collection types which are directly creatable, however, are those directly tagged as such (with the CREATABLE type tag).

In addition to the standard create messages inherited from the defobj Create supertype, Collection inherits the SetIntialValue, Copy, and Drop types (also from defobj). Copying or dropping a collection has no effect on objects which might have be contained in a collection as members; only the references to objects contained in the collection are copied or freed.

SetInitialValue defines special create messages to load the initial members of a newly created collection from an existing collection, and optionally to make the members of a collection read-only after this initialization. If the members of a collection are read-only, all collection messages which add, replace, or remove members will raise an InvalidOperation error. Even if a collection was not initially created with ReadOnly status, this status may also be set later. If a collection is created with the ReadOnly restriction set, however, this restriction cannot later be unset.

A collection used as an initial value need not have the same type as a collection being created. Members (and keys, too, if applicable) are added to the new collection in the same order as defined by traversing the initial value collection with a local temporary index. Create-time options are also applied as noted under SetInitialValue. Some collection subtypes may have further special rules for initial values; if so, these are documented with the subtype.

[aCollection copy: aZone] is equivalent to [[aCollection getType] create: aZone setInitialValue: aCollection].

Create-time options

```
CREATING
- (void) setMemberType: aDataType;
- (void) setIndexSafety: indexSafety;
- (void) setIndexUpdateHandler: (handler_t)handler : (void *)arg;
SETTING
- (void) setReplaceOnly: (BOOL)replaceOnly;
USING
- getMemberType;
- getIndexSafety;
- (handler_t) getIndexUpdateHandler: (void *)arg;
- (BOOL) getReplaceOnly;
```

MemberType

(.. This option is not currently implemented for any collection types.)

The MemberType option may be used to declare the type of member which a collection contains. Its value must be an objec having one of the ValueType types defined in defobj. (..Currently no ValueType objects are implemented, so MemberType is not supported.)

The member type defaults to the Symbol value Unknown. Specifying a member type has no effect on the collection if the storage size of the data type is no larger than a pointer of id type. The member type, however, may be helpful as documentation, and members can be checked for conformance to the type as part of debug options. If the data type has a size greater than id, the collection automatically allocates and frees space for the member value whenever member values are added or removed. For such a collection, all member values must be passed to the collection not as the value itself, but as a pointer to the value (cast to id as necessary). If a new member is being added, the value will be copied to a new member location contained within the collection. Member values returned by the collection will be returned as pointers to this internal value, not as the value itself. In the case of member values removed from the collection, a nil pointer is returned.

IndexSafety

(.. This option is not currently implemented for any collection types.)

The IndexSafety option has one of the following values:

```
id <Symbol>  Unsafe, UnsafeAtMember, SafeAlways;
```

If not specified, the default value depends on the specific collection subtype. In typical cases, the default value is Unsafe. A non-default value is required for indexes to have a guarantee of safety when independent modifications to the underlying collection are also performed.

Without a guarantee of index safety, the effect of any operation which modifies the underlying collection (adding, removing, or replacing members) could be to disrupt the position assumed by any existing index. The effect of such disruption is, in general, completely unpredictable: it could be an immediate crash, a wrong sequence of members, or other difficult-to-debug conditions. To avoid such unpredictability, several means for assuring the safety of indexes under independent collection update are provided.

The simplest way to guarantee index safety is to set the IndexSafety option to the value SafeAlways when a collection is created. With the value of SafeAlways, any active index on the collection will be updated automatically to reflect any change to the underlying collection whenever such a change occurs. This automatic update continues to occur throughout the lifetime of any index on the collection; it ceases to occur only after an index is dropped. This is the easiest option to use, but consumes extra resources to accomplish (both time and space within the collection), so is most suitable when the extra resource usage deesn't matter or there's no better way.

With an IndexSafety value of SafeAlways, there is also an additional option to declare an external index update handler. The handler is specified using the setIndexUpdateHandler:: message. If an IndexUpdateHandler is specified, the value for IndexSafety is also automatically set to SafeAlways.

The first argument of setIndexUpdateHandler:: specifies a function which is called whenever a modification to the underlying collection occurs which could require an update to existing indexes. The second argument establishes a pointer value which the handler function may use to identify the indexes which need to be updated, or any other information needed for the update.

(.. Exact call conventions for the IndexUpdateHandler are still being established, but the activity library will soon be converted to use this technique, so the conventions will be resolved during that conversion.)

Besides the extreme values of Unsafe and SafeAlways, another valid value for IndexSafety is UnsafeAtMember. This value indicates that the only modification of the collection which could disrupt existing indexes is to remove the member contained at the current index position, or to change the member on either side of an index positioned between members. For some collection subtypes, this level of safety is automatically assured, and for others it can be requested with less overall resource impact than SafeAlways.

ReplaceOnly

(.. This option is not currently implemented for any collection types.)

This boolean-valued option restricts valid usage of a collection by excluding all operations which add or remove members. For some collection subtypes, a replace-only restriction can obtain many of the same performance advantages as a read-only collection, but without disabling replace operations as well. Just like the ReadOnly option, the ReplaceOnly option may be reset after a collection is created, provided it was not originally set to true.

Collection behavior

```
- (int) getCount;
```

getCount returns the integer number of members currently contained in the collection. All collections maintain their count internally so that no traversal of collection members is required simply to return this value.

```
- atOffset: (int)offset;
- atOffset: (int)offset put: anObject;

- getFirst;
- getLast;
```

An offset is an integer value that gives relative position of a member in the enumeration sequence of a collection. Offsets start the count of the first member at zero, just like C array indexing.

Offsets provide an alternate means to access the members of a collection, without creating a separate index object. Some collection subtypes (such as Array) support fast, direct access by integer member offset, while others support member offsets only as a shorthand for sequential access through every preceding member. Access by offsets is supported on all collections regardless of whether its speed on a particular collection type.

The atOffset: message returns the member at a particular member offset. The atOffset: put: message replaces the member at a particular offset with a new value, and returns the previous member value at this offset. Both these messages raise the error OffsetOutOfRange if the offset is less than zero or greater than or equal to the count of members in the collection.

The getFirst message is equivalent to [aCollection atOffset: 0], and the getLast message is equivalent to [aCollection atOffset: [aCollection getCount] - 1].

```
- (BOOL) contains: aMember;
- remove: aMember;
- (void) removeAll;
```

The contains: message returns true if the collection contains any member value which matches the value passed as its argument. Depending on the collection subtype, this may require traversing sequentially through all members of the collection until a matching member is found. For other subtypes, some form of direct indexing from the member value may be supported. The message is supported regardless of its speed.

The remove: message removes the first member in the collection with a value matching the value passed as its argument. If there is no such member, a nil value is returned. As with the contains: message, the speed of this operation may vary from very low to linear in the number of members, depending on the collection subtype.

The removeAll message removes all existing members of a collection and sets its member count to zero. The collection then remains valid for further members to be added. This message has no effect on the objects which might be referenced by any removed member values. If resources consumed by these objects also need to be released, such release operations (such as drop messages) can be performed prior to removing the member values.

```
- begin: aZone;
```

The begin: message is the standard method for creating a new index for traversing the elements of a collection. All further information about indexes is documented under the Index type.

```
- (void) forEach: (SEL)aSelector;
- (void) forEach: (SEL)aSelector : arg1;
- (void) forEach: (SEL)aSelector : arg1 : arg2;
- (void) forEach: (SEL)aSelector : arg1 : arg2 : arg3;
```

The forEach messages supply a convenient shorthand for repeatedly performing the same message on all objects contained as members in a collection. The message to be sent is identified by the argument aSelector. This selector must define the same number of arguments as contained in any remaining argument slots of the forEach message. The argument types of the message to be sent must be either the id type, or some other type that will fit in the same number of bits as the id type. By global portability assumptions, the argument type could be as large as an int (signed or unsigned), but not necessarily as large as a long. To use the message, any non-id value must be cast to the id type as part of the call expression.

The forEach: messages are implemented by a simple loop through all members of a collection, using an internal, temporary index. If any operation more complex than a simple message send is required, this operation should just be coded directly using a loop that traverses its own index.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Index – reference into an enumeration sequence for a collection

Synopsis

```
//
// Index -- reference into an enumeration sequence for a collection
//
@deftype Index <DefinedObject, Copy, Drop>- getCollection;

- next;
- prev;
- findNext: anObject;
- findPrev: anObject;

- get;
- put: anObject;
- remove;

- getLoc;
- (void) setLoc: locSymbol;
- (int) getOffset;
- setOffset: (int)offset;

- (int) compare: anIndex;
@end
```

Description

An index is a reference into an enumeration sequence of a collection. Such an enumeration sequence contains all members of the collection in some order. This order will always be consistent with ordering of members in the collection, assuming there is such an ordering. Otherwise, the sequence will still contain all members in some order that remains fixed provided that new members are not added or removed from the collection.

An index is created by a begin: or createIndex: message against a collection. Each major collection type has its own corresponding index type, which supports specialized types of processing against the valid contents of that kind of collection. Once created, an index is a separate object from the collection itself, but it remains valid only so long as the collection itself still exists. Multiple indexes may exist at the same time against the same collection, and each index maintains its own position within an enumeration sequence for the collection.

Many indexes provde the ability modify the collection they refer to, in addition to simply traversing members. An index often provides the complete means for maintaining the contents of a collection, more than could otherwise be performed on the collection itself. The position or other status of the index is automatically updated to reflect any changes made through the index itself.

If changes to a collection are made while other indexes exist, those other indexes could be affected in potentially catastrophic ways. Each collection type documents which kinds of changes can be made without affecting or invalidating existing indexes. The IndexSafety option of Collection provides additional ways to protect indexes against possible effects of independent updates.

Inherited behavior

```
@deftype Index <DefinedObject, Copy, Drop>
```

Each index is a stand-alone object allocated within a zone passed as an argument in the message that created it. This zone need not match the zone of a collection. It is common for index lifetimes to be shorter than their collection. For example, indexes can be created in a temporary scratch zone for use only within a local loop.

Because messages to a collection are the only valid way to create an index, create messages and create-time options are not used with index types. All valid processing on an index is determined by characteristics of the collection from which it is created. Index types are typically named after the type of collection they are created from, and serve principally to define the specific messages valid for an index on that type of collection.

Index objects support the universal messages of the top-level DefinedObject supertype, along with the standard drop: and getZone messages. Even though they cannot be created except from a collection, new index objects can be created from an existing index using the standard copy: message. Each copy refers to the same collection as the initial index, and starts at the same position in its enumeration sequence. In all other respects, however, the new copy is an independent index that maintains its own position under any further processing.

Index behavior

```
- getCollection;
```

getCollection returns the collection referred to by an index. This collection never changes during the lifetime of the index.

```
- next;
- prev;
```

The next and prev messages are the basic means to traverse through the members of a collection.

The next message positions the index to the next valid member after its current position, or to a special position after the end of all valid members. The prev message works similarly, but positions to a valid member preceding the current position, or to a special position preceding all valid members. In addition to repositioning the index, both messages return the new member value to which they are positioned, or nil if there is no such member.

Following is a standard form of loop which may be used to enumerate the members of a collection using an index:

```
index = [aCollection begin: aZone];
  while ( (member = [index next]) ) {
    // do something with member ...
  }
```

When first created, the current position of an index is always set to the special position preceding all valid members. From this initial position, this loop accesses each succeeding member in turn. It assigns each member value to a local variable ("member" in this example) until a nil value is returned.

```
- findNext: anObject;
- findPrev: anObject;
```

findNext: and findPrev: repeatedly perform next or prev until the member value of the index matches their argument. nil is returned if the index reaches the end of valid members without matching the argument. They are only convenience shorthands for loops which could otherwise be programmed using next or prev.

```
- get;
- put: anObject;
- remove;
```

get returns the member value at which the index is currently positioned, or nil if the index is not positioned at a member The put: message replaces the member value at the current index position with its argument. Neither of these messages changes the current position of the index.

The get message provides an alternate way to obtain the current member value in a loop that traverses a collection; its return value is the same as next or prev would return when first positioning to a new member.

The remove message removes the member at the current location of an index, and returns the member value removed. The index position is set to a special position between the members which previously preceded and followed the removed member. If there is no preceding or following member, the index is set to the special location before the start or after the end of all members. After a current member is removed, there is no member at the current index location, but a subsequent next or prev message will continue with the same member that would have been accessed had the current member not been removed.

Both put: and remove raise the InvalidIndexLoc error if the index is not positioned at a current member.

```
- getLoc;
- (void) setLoc: locSymbol;
```

The getLoc message returns a symbol constant which indicates the type of location at which an index is currently positioned. This index location symbol has one of the following values:

```
extern id <Symbol>  Start, End, Between, Member;
```

The Start symbol indicates the special position preceding all members in the enumeration sequence for the collection. This is the location at which an index is positioned when it is first created. The End symbol indicates the special position following all members in the collection. This is the location at which an index is positioned just after a next message has returned nil, as a result of moving beyond the last member. The Between symbol indicates the special position between two other members, at which an index is positioned after a current member, between two other members, is removed. The Member symbol indicates that the index is positioned at some current member in the enumeration sequence of a collection.

The getLoc message is needed to traverse a collection which could contain nil values for its members. Without getLoc, there would be no way to distinguish a nil value returned by next as either a valid member value or the special value returned at the end of members. With getLoc, a loop that traverses a collection can test specifically for the end (or start) of members. Following is a simple loop which illustrates such usage: index = [aCollection begin: aZone]; do { member = [index next]; // do something with member ... } while ([index getLoc] == Member); The setLoc: message may be used to reset the current location of an index to either Start or End. It may be used to reprocess a collection using an existing index after some other location has already been reached. It may also be used to position an index at the end of all members prior to traversing members in reverse order using prev.

Besides Start and End, setLoc: accepts the special argument values of BetweenAfter and BetweenBefore, which are also defined symbols. These argument values are only valid if the index is positioned at a current member. They reposition the index to the special location between the current member and its immediately following or preceding member.

```
- (int) getOffset;
- setOffset: (int)offset;
```

Provided there is no major computational cost, an index also maintains the integer offset of its current member within the enumeration sequence of the collection. These integer offset values have the same definition as in the atOffset: messages of Collection. The getOffset message returns this current offset value. If the index is current positioned at the Start or End location, getOffset returns -1. If the index is positioned at a Between location, getOffset returns the offset of the immediately preceding member. If the offset is not currently available from the index, getOffset returns the special value UnknownOffset. This value is defined by a macro as the maximally negative value of a 32-bit, 2's-complement integer.

An offset is always available from an index if its current position has been reached by repeated operations from an initial Start or End position, and there has been no other modification to the underlying collection. Some forms of direct member access operations supported by some index types, however, may result in an integer offset not being available. These restrictions are noted with the individual index type.

Using the setOffset: message, an index may be positioned directly to a member using the offset of the member within its enumeration sequence. The speed of this operation depends on the specific type of the collection, just as noted for the atOffset: message on Collection. In the worst case, this operation is linear in the magnitude of the offset.

```
- (int) compare: anIndex;
```

The compare: message compares the current location of one index with the current location of another index passed as its argument. If the two indexes have the same location, compare: returns zero. Otherwise, compare: returns +1 or -1 according to whether the argument index precedes or follows the receiver index in the enumeration sequence of the collection. If either of the two indexes has an unknown offset, and the location of the other index is anything other than Start or End or an immediately adjacent member, compare: returns the UnknownOffset integer value.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Array – collection supporting access to fixed group of members

Synopsis

```
//
// Array -- collection supporting access to fixed group of members
//
@deftype Array <Collection, CREATABLE>CREATING
+ create: aZone setCount: (int)count;
+ create: aZone setMemberBlock:(id *)members setCount:(int)count;
SETTING
- (void) setCount: (int)count;
- (void) setDefaultMember: memberValue;

- (void) setMemberBlock: (id *)members setCount: (int)count;
USING
- getDefaultMember;
- (id *) getMemberBlock;
@end
```

Description

An array is a collection of members that are all created as members of the collection at the same time. Existing member values may be replaced with new values, but the members themselves are fixed at constant offsets within the collection. The fixed structure of an array permits very fast access to members by integer offset positions, since the location of each member may be directly calculated.

The Array type is one of the simplest collection types in the collections library, and the closest to a data structure directly supported in C. Unlike C arrays, the group of members belonging to the array is not necessarily fixed for the lifetime of the array, but may be dynamically resized to contain a different number of members. When an array is dynamically resized, existing member values are preserved as much as possible.

The Array type adds few messages to the generic messages inherited from Collection. This type is provided partly so that a fixed-structure array can be accessed with the same uniform set of basic messages as any other kind of object collection. It also handles all required memory allocation within the collection. As an option, however, the Array type can be used to wrap an existing C array for external access as an object collection. It can also provide access to an internal C array for direct manipulation using C expressions. These forms of low-level access support hybrid modes of use in which advantages of both low-level manipulation and external object access can be combined.

Inherited behavior

```
@deftype Array <Collection, CREATABLE>
```

The Array type is directly creatable, and supports all standard messages of Collection except removal of individual members. The messages based on an integer offset, either on the collection (atOffset:, atOffset:put:), or an index (setOffset:) all execute in fast constant time. Members of an array are fully ordered according to these integer offsets. Sequential access to members through its members is also fully supported. The Array type disables the remove message inherited from Collection; the message is defined, but any attempt to remove a member will raise an error that the operation is not supported.

While the IndexSafety option of an Array defaults to Unsafe, indexes are inherently safe under all operations with the exception of resizing an array or resetting its member allocation. (The only individual member operation which modifies contents of an array is replacing a member value, and this does not affect existing indexes.) The UnsafeAtMember value of IndexSafety is not supported since individual members cannot be added or removed. A value of SafeAlways makes existing indexes safe even under resizing or reallocation. If an existing index referred to a member location that is no longer contained in the array, the index location is reset to the End location.

The default value of the ReplaceOnly option is true, and cannot be overridden.

The type of index returned by begin: on an array is simply Index. There is no special index type for Array because there are no additional messages beyond those already defined by Index.

Create-time options

```
SETIING
- (void) setCount: (int)count;
- (void) setDefaultMember: memberValue;

- (void) setMemberBlock: (void *)members setCount: (int)count;
USING
- (void) setDefaultMember: memberValue;
- getDefaultMember;

- (void) setMemberBlock: (void *)members setCount: (int)count;
```

```
- (void *) getMemberBlock;
```

All the Array create-time options can also be set after the array is already created, subject to restrictions noted below.

Count

The Count option sets the number of members which belong to the collection. Any non-negative value including zero is valid. If the array already exists, the any existing members up to the new count will preserve their existing values. If the new count is greater than the existing count, or a new array is being created, all members will be assigned an initial default value of either nil, or a value previously specified for DefaultMember.

The inherited getCount message of Collection obtains the current value set for this option. There is no way to add or remove members of an array except using the setCount: message directly.

DefaultMember

When this option is set, the initial value of all new members will be set to the member value given (otherwise the default is nil). This option gives a convenient way to distinguish members which have never been set from any other valid member value, which could include nil. This option may be reset after array creation only if some setting for the option was given at create time. (The initial, explicitly set value can still be the default nil, but a value must be set explicitly for the option to be resettable later). The get message for this option always retrieves the current setting, but this value has no effect except when the count of an array is increased, so that new members need to be initialized.

MemberBlock

This option provides a means to wrap an existing C array for access as an object collection. If this option is given, dynamic resizing is not supported. The current C array being wrapped, however, can be replaced by giving a new setting for MemberBlock after an array already exists. A new setting can be given only if some setting was initially given at create time.

Even if a setting was not given for MemberBlock at create time, the get message for MemberBlock always returns a pointer to whatever internal memory array is currently being used by the array object. If a MemberBlock setting was given, the pointer returned will be the same as the one previously given. In either case, the pointer returned may be used to manipulate member values in any way desired using native C expressions. After an array has been created, there is no way to determine whether the MemberBlock pointer was established externally or by internal allocation. If the pointer was established by internal allocation, however, the external program must make no attempt to free or otherwise reallocate this memory.

A count must always be supplied with an external member allocation, using the setCount: argument of the compound message. If an external allocation is being used, the only way to reset the count is also to reset MemberBlock; any attempt to use the setCount: message by itself will raise an error. Whenever an external member allocation is being used, the external program is entirely responsible for assuring that the MemberBlock value is a pointer to valid allocated memory containing at least the number of member slots given by setCount:.

With an external member allocation, the array itself will not attempt to either allocate this memory or free it when the array is dropped. Dropping the array only removes its reference to the external allocation.

The pointer to allocated memory is given as a pointer to id type, but if a data type given as MemberType has size greater than id, the pointer must actually point to an array of slots each of which has the size of the data type.

Since an array neither allocates nor frees an external member allocation, the same region of allocated memory may be referenced by multiple arrays, including overlapping member ranges each defined by starting location and count. This flexibility enables alternate subrange views of a single, contiguous initial allocation by means of separately created external collections.

Create combination messages

```
CREATING
+ create: aZone setCount: (int)count;
+ create: aZone setMemberBlock:(id *)members setCount:(int)count;
```

In addition to the standard inherited create messages, the Array type supports two forms of combination messages which both create the array and set its initial options. These are provided both for convenience and to provide the fastest possible array creation for simple cases in which speed might be important. The two cases supported are creating an array with a member count for its internal members already set, and creating an array which refers to an external member allocation. If any more options than these are needed, the normal method of setting options between createBegin: and createEnd messages should be used instead.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

List, ListIndex – collection of members in linear sequence

Synopsis

```
//
// List - collection of members in linear sequence
//
@deftype List <Collection, CREATABLE>CREATING
- (void) setEndsOnly: (BOOL)endsOnly;
USING
- (BOOL) getEndsOnly;

- (void) addFirst: anObject;
- (void) addLast: anObject;

- removeFirst;
- removeLast;
@end

//
// ListIndex -- index with insertion capability at any point in list
//
@deftype ListIndex <Index>- (void) addAfter: anObject;
- (void) addBefore: anObject;
@end
```

Description

A list is a collection of members that are all maintained at some externally assigned position in a linear sequence of all members. The sequence is established by the position at which members are added: members can be added at the start of list, at the end, or at any point in the middle.

A list is also one of the most dynamic of basic collections, in that the cost of adding and removing members is very low. Members can be removed from any position just as easily as they can be added. A list automatically grows and shrinks to reflect the number of members at any one time, and there is no fixed capacity which limits the size to which a list may grow.

Inherited behavior

```
@deftype List <Collection, CREATABLE>
```

The List type is supports all messages of Collection. If created with default options, it provides no special speedup of accesses by integer offset.

Unless the EndsOnly option is specified, the default value of IndexSafety is UnsafeAtMember. This means that indexes are inherently safe under all operations except those which affect the member at an index or immediately adjacent to an index between members. If the EndsOnly option is specified, the default value of IndexSafety is Unsafe. In this case, indexes could be invalidated by any operation that adds or removes members. The SafeAlways can be specified for any case to get a guarantee of full index safety.

ListIndex is a special index type just for list collections. It provides the only method to add members within the interior of a list.

Create-time options

```
CREATING
- (void) setEndsOnly: (BOOL)endsOnly;
USING
- (BOOL) getEndsOnly;
```

EndsOnly

(.. This option is not implemented yet.) The EndsOnly option restricts a list so that members can be added or removed only at either end. This restriction permits an alternate implementation that both reduces memory usage and enables fast, direct access to members by integer offset. It obtains key performance advantages of an array while preserving the dynamic extendability of a list.

List behavior

```
- (void) addFirst: anObject;
- (void) addLast: anObject;
```

The addFirst: and addLast: messages add a new member to either end of a list. With addFirst:, the new member becomes the first member of the array. With addLast:, it becomes the last member.

```
- removeFirst;
- removeLast;
```

removeFirst and removeLast remove either the first or last member of a list. They each return the value of the member just removed.

ListIndex behavior

```
//
// ListIndex -- index with insertion capability at any point in list
//
@deftype ListIndex <Index>- (void) addAfter: anObject;
- (void) addBefore: anObject;
@end
```

The addAfter: and addBefore: messages add members at a particular point in the sequence of members maintained by a list. The current location of an index determines the point at which a new member will be added. The addAfter: message adds a member at the list position immediately following the current index location. addBefore: adds a member to the immediately preceding location. Neither message changes the current location of the index, except that an index can change from a Start or End location to a location of Between.

Since an index may be positioned to any location in a list, these messages enable the construction of any desired sequence of members. Since the current index location remains unchanged, multiple members may all be inserted successively at some point in a list; previously added members are just pushed out one-by-one as new members are added.

An index with a location of Start, End, or Between is just as valid a location for addAfter: or addBefore: as an index positioned at a member. In these cases, there is no member at the current location of the index, so the new member is just inserted directly at the current index location, and the index is left positioned between the new member and the member that was previously adjacent in the opposite direction. If the previous location was Start and the message addAfter:, or the location was End and the message addBefore:, the index location remains Start or End.

If either the addAfter: or addBefore: message is requested with the EndsOnly option, and the index location is not Start or End, or if remove is requested on an index that is not positioned at either the first or last member, an invalid operation error is raised.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Stack, Queue – restricted forms of lists

```
(..  Neither of these types is currently implemented.  They are nothing more than
restricted cases of a List, so if you need them use a List instead.)
```

Synopsis

```
// Stack -- form of list with add/remove both at start of list
//
@deftype Stack <List>- (void) push: anObject;
- pop;
@end

//
// Queue -- form of list with add at start of list and remove at last
//
@deftype Queue <List>- (void) enqueue: anObject;
- dequeue;
@end
```

Description

```
A Stack is a subtype of List which restricts the addition and removal of members to the start
of a list.  Queue is a subtype of List which restricts addition of members to the start of
the list and removal to the end.  Both types define specially named messages to reflect the
specialized usage of their resulting add and remove operations.
```

Interface Reference

Inherited behavior

```
        @deftype Stack <List>@deftype Queue <List>
...
```

Stack behavior

```
- (void) push: anObject;
- pop;
...
```

Queue behavior

```
- (void) enqueue: anObject;
- dequeue;
...
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

KeyedCollection, KeyedIndex – supertype of Set and Map types

Synopsis

```
//
// KeyedCollection -- supertype of Set and Map types
//
@deftype KeyedCollection <Collection>CREATING
- (void) setDuplicatesOption: duplicatesOption;
- (void) setGroupType: aCollectionType;
- (void) setSingletonGroups: (BOOL)singletonGroups;

- (void) setSorted: (BOOL)sorted;
- (void) setCompareFunction: (idXid_t)aFunction;
-
(void)        setBucketFunction: (bucket_t)aFunction setBucketCount: (int)n;

- (void) setPartiallyOrdered: (BOOL)partiallyOrdered;
- (void) setMemberSlot: (int)byteOffset;
SETTING
- (void) setPartialOrderContext: aKeyedCollectionCollection;
USING
- getDupOption;
- getGroupType;
- (BOOL) getSingletonGroups;

- (BOOL) getSorted;
- (int(*)()) getCompareFunction;
- (int(*)()) getBucketFunction;
- (int) getMaxBuckets;

- (BOOL) getPartiallyOrdered;
- getPartialOrderContext;

- (int) getMemberSlot;

- at: aKey;
- (BOOL) containsKey: aKey;
- (int) getCountAtKey: aKey;

- removeKey: aKey;
```

```
- createIndex: aZone at: aKey;
- createIndex: aZone setMember: aMember;

- (void) getPredecessors: aMember;
- (void) getSuccessors: aMember;
- (void) getCostarts: aMember;
- (void) getCoends: aMember;
@end

//
// KeyedIndex -- index behavior shared by Set and Map types
//
@deftype KeyedIndex <Index>
- (int) getCurrentLevel;
- getIndexAtLevel: (int)level;

- at: aKey;

- setMember: aMember;
@end
```

Description

A keyed collection is a collection in which each member can be compared with some other value that identifies the member. This value is referred to as the member key. The key value may be determined either by the member value itself, which defines a Set, or by external association with the member when the member is first added, which defines a Map.

The keyed collection type establishes the common behavior shared by both Set and Map. Standard options are provided to declare ordering of members in the collection, which can be based either on an inherent ordering of the key values (a sorted collection) or external maintenance of a total or partial order. Additionally, options are provided to declare whether duplicate key values should be permitted, and if so, how.

The options for handling of duplicate keys provide considerable flexibility. Entire structures of nested collections can be created that subdivide the key space through progressive levels. By accepting duplicate keys, and organizing members of the same identity into subgroups, keyed collections support the abstract data type known as a bag or multiset. The index on a keyed collection includes the ability to traverse throughout the contents of any of these nested collections.

Inherited behavior

```
@deftype KeyedCollection <Collection>
```

The KeyedCollection type inherits all standard behavior of Collection. The KeyedCollection type is not itself creatable; it only serves as a common supertype for Set and Map collection types.

The default value of IndexSafety for a KeyedCollection is Unsafe, but values of either UnsafeAtMember or SafeAlways can be requested to guarantee a higher level of safety.

Create-time options

(.. None of these options is implemented yet, except for CompareFunction.)

```
CREATING
- (void) setDupOption: dupOption;
- (void) setGroupType: aCollectionType;
- (void) setSingletonGroups: singletonGroups;

- (void) setSorted: (BOOL)sorted;
- (void) setCompareFunction: (idXid_t)aFunction;

- (void) setPartiallyOrdered: (BOOL)partiallyOrdered;

- (void) setMemberSlot: (int)byteOffset;
SETTING
-
(void)        setBucketFunction: (bucket_t)aFunction setBucketCount: (int)n;
- (void) setPartialOrderContext: aKeyedCollection;
USING
- getDupOption;
- getGroupType;
- (BOOL) getSingletonGroups;

- (BOOL) getSorted;
- (int(*)()) getCompareFunction;

- (BOOL) getPartiallyOrdered;
- getPartialOrderContext;

- (int) getMemberSlot;

- (int(*)()) getBucketFunction;
- (int) getBucketCount;
```

DupOption

The value of DupOption specifies how duplicate key values are to be handled within a keyed collection. This option has one of the following values:

```
id <Symbol>  DupIsError, DupRejected, KeepDuplicates, KeepCountOnly;
```

If the DupIsError value is specified, any attempt to insert a duplicate key is an error, and effects of such an attempt depend on the level of debug checking. If debug checking for this error is turned on, the collection will check whether any new key duplicates an existing one, and raise a DuplicateKey error if it does. If debug checking is turned off, there is no guarantee that any check will be made, and the effects could be unpredictable in or catastrophic or insidious ways.

With a value of DupRejected, attempts to insert a duplicate key are always detected, and the attempt rejected. (All insert messages return a boolean flag as to whether a duplicate occurred.)

With a value of KeepDuplicates, attempts to insert a duplicate key are always accepted, and all members having the same key are retained in a specially created internal collection containing just the members that all have that same key. This collection then becomes the single member internally associated with the key value.

With a value of KeepCountOnly, duplicate members are not retained. Instead, only a count of the net number of insertions and removals at that key value are retained within the collection. With this value, member values cannot be retrieved from the collection, but only a count of members for a given key value, using different messages. This value is accepted only for the Set subtype of a KeyedCollection.

GroupType

The GroupType option specifies the type of internal collection to be created for members that all have the same key. This option is valid only if DupOption is also set to a value of KeepDuplicates.

Without this option, the type of internal collection created for members with the same key depends on the specific subtype. This option allows an external program to customize the type of internal collection permitted. The argument must be a previously created subtype of Collection. Typically, the argument is a customized version of one of the standard collection types such as List, Set, or Map; the argument would be the type object returned by customizeEnd after a customizeBegin: sequence.

More information about all the ways this option can be interpreted is provided with each subtype.

SingletonGroups

If this option is true, the type of collection created for members with duplicate keys is also created even when there is only member for a particular key. This option can be convenient to assure uniformity of a keyed collection with nested structure. This option is valid only if DupOption is also set to a value of KeepDuplicates.

Sorted

If this option is true, the immediate members of a keyed collection (included collections of duplicate members, if any) are totally ordered according to an ordering relation defined on key values. The default order for sorting is determined by a standard compare: message on a key value to be added vs. existing keys of the collection. This default method to establish order can be overridden using the CompareFunction option.

The default for the Sorted function is true, unless a value is also given for the option BucketFunction.

CompareFunction

Use the function pointed to by the argument as the method for determining whether two members have the same key value, and also for ordering the keys of the collection if the Sorted option is true. The argument is declared with the following function pointer type:

```
typedef int (*idXid_t)( id, id );
```

The function given will be called whenever one key value needs to be compared with another. Multiple calls to the function typically occur whenever members are added or removed from the collection, until the correct member for insertion or removal is determined.

The compare function accepts two key values as its arguments. Like member types, key types declared with the id pointer type, but may contain any type of value (such as an integer) up to the size of an id. An explicit compare function can support any type of value as a key, regardless of whether it is an object that supports a standard compare: message.

The compare function is called repeatedly by the collection to compare two key values. The function should return zero if the key values are equal, -1 for the first key argument less than the second, and +1 for the first greater than the second. If a keyed collection is not sorted, either -1 or +1 may be returned for unequal keys, regardless of whether one might be taken as greater or less than the other.

BucketFunction

The BucketFunction option specifies a function that is called whenever a member of an unsorted collection is added or removed. It may also be used for a sorted collection to provide faster insertion and removal of members. The responsibility of a "bucket function" is to assign the key value to some integer in a range from zero up to some maximum value. These integers identify "buckets" to which the same key value will also be assigned whenever passed to the function. If the collection is unsorted, these buckets represent hash buckets, and for speed should be relatively sparse (few keys likely in any one bucket, given the expected size of the collection). If the collection is sorted, these buckets represent ordered buckets to which the same range of key values will always be assigned. Such sorted buckets can speed up the search for a key location within a sorted collection. Sorted buckets can also be sparse to minimize the need for further sorting within each bucket. If many keys end up in the same sorted bucket, the sorting process normally used without buckets is still used, so there is less performance impact than with hash keys of an unsorted collection.

A bucket function accepts a single key value as its argument, and return an integer that identifies a bucket. As with a compare function, the key value is declared as an id pointer type, but may actually contain whatever kind of value is used within the collection. Following is the type of the bucket function pointer:

```
typedef int (*bucket_t)( id );
```

A new bucket function may be established any time during the lifetime of a collection. A new bucket function can ensure that keys remain suitably distributed to buckets as the total number of keys in a collection grows or shrinks.

BucketCount

A value for MaxBuckets must be given in the same message that sets a bucket function. Its value specifies the number of buckets to which a bucket function assigns its keys. The maximum value returned by a bucket function should be one less than this value.

PartiallyOrdered

A partially ordered collection supports the maintenance of individual ordering relations among its members. Each such relation is an assertion that one member should always precede another in the enumeration sequence of the collection. When an index traverses the collection, it will guarantee that all the predecessors of a member have already been visited before the member itself is reached.

If true, PartiallyOrdered specifies that messages for maintaining a partial order are enabled on the collection. An error will be raised if individual members do not have a unique identity within the innermost collection within which they are contained. This means that if duplicate keys are permitted by the collection, a collection type must also be specified for GroupType which does not itself accept duplicates. Every member in a partial order must have a unique identity so that the member itself is sufficient to identify it within the network of partial orders.

PartialOrderContext: aKeyedCollectionCollection

Ordering relations may be maintained not only between members belonging to the same collection, but between members and other collections, as long as all participants in an ordering are contained somewhere within a larger common collection. An ordering relation with a nested collection specifies that all members in the collection are successors or predecessors of the other member. Ordering relations can also be established between members of a larger, containing collection, and individual members of a nested collection.

The keyed collection specified as the value for PartialOrderContext must specify another collection in which the local collection is maintained. It is an error if the PartialOrderContext collection does not actually contain the local collection.

getMemberSlot

The MemberSlot option indicates that space has been reserved within each member that allows it to contain a link directly to its position in the enumeration for a collection. If such space has been reserved, special messages can be used that rapidly position an index directly to the member. Operations to remove members are also much faster.

The value of MemberSlot specifies the offset in bytes from the start of each member where the space for its position link has been reserved. The following typedefs are provided to declare the space required for a position link within a member allocation:

```
typedef struct memberData { void *memberData[2]; }            member_t;
typedef struct memberData { void *memberData[2]; id owner; }  dupmember_t;
```

The first type declares the required space for a position link in a collection that does not accept duplicate key values. The second type is required if duplicate keys are permitted.

The offset of the position link from the start of a member may be either positive or negative, in the range of -2048 to +2047. The default value of MemberSlot is UnknownOffset (a large negative value), which specifies that no slot for an internal position link is available within each member.

If a member slot is defined, a memory pointer of some kind must always be used for the value of every member. These members may be object id pointers, but other types of memory pointers are acceptable as well.

The contents of the first two words of position link vary according to the position of a member in a collection. The third word of a dupmember_t link contains the id of the collection in which a member is directly contained. This collection is either an internally created collection containing members with duplicate keys, or the collection to which the member was added if there is no other member with the same key.

KeyedCollection behavior

(.. None of the messages or other behavior relating to duplicate key values are currently implemented.)

```
- at: aKey;
- (BOOL) containsKey: aKey;
- (int) getCountAtKey: aKey;  // unimplemented

- removeKey: aKey;
```

The at: message returns the existing member of the collection which matches the key value passed as its argument, or nil if there is no key value in the collection which matches. If duplicate entries for this key exist, the entire collection of duplicate members created for the key value is returned instead.

The containsKey: message returns true if the key value passed as its argument is contained in the collection, and false otherwise. The getCountAtKey message returns the number of member entries which are all associated with its key value argument. This number is zero if the key value is not contained in the collection, and will never exceed one one unless a collection accepts duplicate entries for a key.

The removeKey: message removes a member matching a key value from the collection, and returns the member just removed. It returns nil if there is no key value in the collection which matches. If more than one entry was present for the key value, it removes and returns the first member in the internal collection created for duplicate members.

```
- createIndex: aZone at: aKey;
- createIndex: aZone setMember: aMember;
```

(.. Neither of these messages is currently implemented.)

The createIndex:at: message creates a new index on the collection and also sets its current position to the first member matching the key value given as its final argument. It returns the index just created.

The createIndex:at: message creates a new index on the collection and also sets its current position to the member passed as its final argument. This message is valid only if an internal member slot was defined for the collection with the MemberSlot option.

```
- (void) getPredecessors: aMember;
- (void) getSuccessors: aMember;
- (void) getCostarts: aMember;
- (void) getCoends: aMember;
```

(.. None of these messages is currently implemented.) These messages maintain the ordering relations within a partially ordered collection.

KeyedIndex behavior

```
- (int) getCurrentLevel;
- getIndexAtLevel: (int)level;
```

An index to a keyed collection traverses all members of the collection, regardless of whether these members belong to collections of members entered under duplicate key values. Internally, however, an index keeps track of any specific subcollection it is currently processing. The getCurrentLevel message returns the level of nesting, within collections of duplicate entries for a current key value, to which the current position of an index belongs. The level returned is an integer that starts at zero for the top-level collection, and increases by one for each level of nested collection that is currently active.

While processing any nested collections, the index for a keyed collection actually maintains a separate, internal index for each level of collection that is currently active. The index at each level may be retrieved using the getIndexAtLevel: message.

```
- at: aKey;
```

The at: message repositions the index to the first member that matches the key value passed as its argument.

```
- setMember: aMember;
```

(.. This message is not currently implemented.) The setMember: message repositions the index to the member passed as its argument. This message is valid only if an internal member slot was defined for the collection with the MemberSlot option.

Roger Burkhart <rmb@santafe.edu>
(mailto:rmb@santafe.edu)

Last modified: 1997-12-19 (mgd)

Name

Set – collection of members each having a defined identity

Synopsis

```
//
// Set -- collection of members each having a defined identity
//
@deftype Set <KeyedCollection, CREATABLE>CREATING
- (void) setKeyFunction: (id (*)(id))aFunction;
USING
- (id (*)(id)) getKeyFunction;

- (BOOL) add: anObject;
- (BOOL) add: anObject setIndex: anIndex;
- replace: anObject;
@end
```

Description

Set is a subtype of KeyedCollection in which the key value associated with each member is determined by the member value itself. The key value may be identical to the member itself, or may be defined as a a function of the member using a create-time option.

Inherited behavior

```
@deftype Set <KeyedCollection, CREATABLE>
```

The Set type inherits most of its interface from the KeyedCollection supertype. Set defines no create-time options beyond those already defined by KeyedCollection. If a custom compare or bucket function is specified, the member value is passed as the key value arguments of these functions. These functions determine what part of the member value is part of the key value, by determining which key values will compare equal to any member.

Set behavior

```
- (BOOL) add: anObject;
- (BOOL) add: anObject setIndex: anIndex;
- replace: anObject;
```

```
@end
```

The add: message adds a new member to a set. It returns true if member added matched the key of any member already contained in the Set. Whether or not the member was actually added to the collection or not depends on the setting of DupOption.

The add:setIndex: message adds a new member to a set just like add:, but also has a side effect of setting the position of the index passed as the setIndex: argument to the member which was just added. If an index positioned to the new member is needed anyway, setting the index in the same operation avoids repeating a search for the position at which the member key belongs.

The replace: message replaces a member stored at a given key with another member value that matches the same key. The new value to replace the existing one is passed as the argument. replace: returns the member value that was replaced, or nil if the collection contained no member with a matching key.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

OrderedSet – set of members in an externally assigned linear sequence

```
(..  This type is currently implemented only using the low-level option of an internal
member slot, and the messages for that option do not match the documentation in KeyedCollection.
If you need one of these objects, then either use a List or wait for some other implementation.)
```

Synopsis

```
//
// OrderedSet -- set of members in an externally assigned linear sequence
//
@deftype OrderedSet <KeyedCollection, List, CREATABLE>@end
```

Description

```
An OrderedSet is a totally ordered collection of members in which every member also has a distinct
identity as defined by comparison against a key value.
```

Inherited behavior

```
        @deftype OrderedSet <KeyedCollection, List, CREATABLE>
The sequence of members of an OrderedSet is established using the same messages that maintain
member sequence in a List.  An OrderedSet supports customization and access by key as defined
by Set and KeyedCollection.  The union of messages from all these sources defines the total
interface of an OrderedSet.  Members with duplicate keys, however, are not valid for an OrderedSet.
Each member must have a unique position within the member sequence
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Map – collection of associations from key objects to member objects

Synopsis

```
//
// Map -- collection of associations from key objects to member objects
//
@deftype Map <KeyedCollection, CREATABLE>CREATING
- (void) setKeyType: aDataType;
- (void) setKeySlot: (int)byteOffset;
USING
-          getKeyType;
- (int) getKeySlot;

- (BOOL) at: aKey insert: anObject;
- (BOOL) at: aKey insert: anObject setIndex: anIndex;
- at: aKey replace: anObject;

- (BOOL) insertGroup: aKey;

- createIndex: aZone setKey: aKey;
@end

@deftype MapIndex <KeyedIndex>- getKey;
- setKey: aKey;

- get: (id *)key;
- next: (id *)key;
- prev: (id *)key;

- replaceKey: aKey;
end
```

Description

Map is a subtype of KeyedCollection in which the key value associated with each member is independent of the member itself. Whenever a new member is added to the collection, a key value to be associated with the member must be supplied also. A Map defines a mapping from key values to member values.

Inherited behavior

For the Map type, key values are independent of the member values with which they are associated. Map defines two additional options to document information about its key values. Map also defines its own messages to distinguish the key value from member value in any operation which involves both.

```
@deftype Map <KeyedCollection, CREATABLE>
```

Create-time options

```
CREATING
- (void) setKeyType: aDataType;
- (void) setKeySlot: (int)byteOffset;
USING
-         getKeyType;
- (int) getKeySlot;
```

(..None of these options is currently implemented.)

Map behavior

```
- (BOOL) at: aKey insert: anObject;
- (BOOL) at: aKey insert: anObject setIndex: anIndex;

- at: aKey replace: anObject;
```

at:insert: inserts an entry into a Map containing the key and member values given as its arguments. It returns true if the key was not previously contained in the collection. If the key was already present, the action of the message depends on the value in effect for the DuplicatesOption. With default options, the attempt to insert a duplicate key is simply rejected and false is returned.

at:insert:setIndex: is similar to at:insert:, except that the index passed as the final argument is repositioned to the member just inserted. If further operations are required on the entry in a map that was just inserted, an index provides faster repeated access to this entry than a message involving further key comparison, such as at:, would require. (.. This message is not currently implemented.)

at:replace replaces an existing member value associated with a key value by a new value given as its final argument. The message returns the member value which was formerly associated with the key value. The message returns nil if more than one member is present with the same key.

```
- createIndex: aZone setKey: aKey;
```

(.. This message is not currently implemented. It will eventually give a fast way to create an index already positioned at an entry with a particular key value.)

MapIndex behavior

```
- getKey;
- setKey: aKey;
```

The getKey message returns the key value associated with the current location of the index. It returns nil if the index is not currently positioned at a member. The setKey: messages repositions the index to an entry having a key value that matches its argument. If there is more than one entry matching this key value, the index is positioned to the first entry that matches.

```
- get: (id *)key;
- next: (id *)key;
- prev: (id *)key;
```

These messages are similar to those defined by the Index supertype, except that they combine the return of a member value with an output value for an associated key value. In each case, the key value is passed as an (id *) output argument.

```
- replaceKey: aKey;
```

(.. This message is not currently implemented.)

> replaceKey: replaces a key value with another key value that still compares equal to a key value already present. It may be used to replace an allocation of one object, already contained in a Map, by a different object allocation. replaceKey: returns nil if no existing key value matches its key value argument; otherwise it returns the key value which was replaced by its argument.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

String – character string object

(.. This type will be more fully documented when it is generalized beyond its current status as nothing more than an object wrapper for a null-terminated character array.)

Synopsis

```
//
// String -- character string object
//
@deftype String <Create, Drop, Copy, CREATABLE>CREATING
+ create: aZone setC: (const char *)cstring;
SETTING
- (void) setC: (const char *)cstring;
USING
- (const char *)getC;
- (void) appendC: (const char *)cstring;

- (int) getCount;
- (int) compare: aString;
@end
```

Description

The String object type packages a null-terminated, C-format character string into an object. All memory allocation needed to hold the string value is handled by the object. This type currently defines only the most rudimentary operations for initializing and appending C-format character strings. These are sufficient for its current limited roles in places that need a uniformity between character strings and other kinds of allocated objects.

Inherited behavior

```
        @deftype String <Create, Drop, Copy, CREATABLE>
...
```

Create-time options

```
        CREATING
        + create: aZone setC: (const char *)cstring;
        SETTING
        - (void) setC: (const char *)cstring;
...
```

String behavior

```
        - (const char *)getC;
        - (void) appendC: (const char *)cstring;

        - (int) getCount;
        - (int) compare: aString;
        @end
...
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

OutputStream – stream of output bytes

```
(..  This type will be more fully documented when it is generalized beyond its current status
as a placeholder used only to direct debuggging output.)
```

Synopsis

```
//
// OutputStream -- stream of output bytes
//
@deftype OutputStream <Create, Drop, CREATABLE>CREATING
+ create: aZone setFileStream: (FILE *)fileStream;
- setFileStream: (FILE *)fileStream;
USING
- (FILE *) getFileStream;
- (void) appendC: (const char *)cstring;
@end
```

Description

```
The OutputStream type currently supports only the writing of a character string to a file.
It is a placeholder for more general stream types.  A stream is a collection that supports
only sequential addition of members (an output stream) or sequential removal of members (an
input stream).
```

Inherited behavior

```
        @deftype OutputStream <Create, Drop, CREATABLE>
...
```

Create-time options

```
CREATING
- setFileStream: (FILE *)fileStream;
USING
- (FILE *) getFileStream;
```

```
...
```

Special create: message

```
+ create: aZone setFileStream: (FILE *)fileStream;
...
```

OutputStream behavior

```
- (void) appendC: (const char *)cstring;
@end
...
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Subclassing Reference

Until the collections library has been fully implemented, subclassing conventions from collections implementation classes are still in flux. In general, these classes will be among the most complex uses of multiple classes selected to implement an independent object type. (See library definition conventions (../defobj/library.html) for a summary of the distinction between types and classes.) New methods are being developed to simplify subclassing from such implementations. In the meantime, if you need to use to a collection within the implementation of your own class, just put an instance variable in your class and put the collection in that, and pass through the messages of the collection you want to have available on your class to this variable. In many if not most cases, this is better design anyway, because you control all use of the underlying structure.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Interface Design Notes

Design Objectives

A collections library is one of the most important foundation services for object-oriented programming. Most object-oriented systems provide at least the start of a general-purpose collections library. The GNUSTEP, project, for example, provides the libobject library (currently in alpha test at ftp://alpha.prep.ai.mit.edu/) which includes a collections library along with other services intended to parallel those of the OpenStep framework developed by Next.

Swarm has implemented its own collection library to meet the specialized needs of its agent simulation framework. Major objectives of the Swarm collections library include the following:

... (.. more to come, obviously)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Collections Library: Implementation Notes

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library

Documentation and Implementation Status

Overview

The activity library is responsible for scheduling actions to occur within a simulated world, and for making these actions actually happen at the right time in the right order. It provides the foundation of dynamic, object-oriented simulation within Swarm.

> Actions consist of messages to objects, calls to functions, or groups of actions in some defined order. The activity library guarantees that all these actions, and the state changes they produce, occur at predictable points in time. Time is defined by the relative order of actions, and may also be indexed by the discrete values of a world clock.

Documentation

The activity library follows the Swarm library interface conventions

(defobj/library.html)
and has the following standard sections of documentation:

- Usage Guide
(activity/guide.html)
- Advanced Usage Guide
(activity/adv.html)
- Interface Reference
(activity/ref.html)
- Subclassing Reference
(activity/subclass.html)
- Interface Design Notes

(activity/design.html)

- Implementation Notes
(activity/impl.html)

---

*Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)*

Last modified: 1998-01-08

Activity Library: Documentation and Implementation Status

The activity library is in process of being converted to draw even more of its support from underlying types and classes in the defobj and collections libraries. The basic interface has already been changed to reflect this consolidation, so the programming interface for currently available function is expected to remain stable.

The activity library follows the documentation structure of the Swarm library interface conventions
(../defobj/library.html)
. There are placeholders for each section of documentation so that all links should at least link up with something, whether or not there's anything there. The Interface Reference section is already fairly complete, and there is a start toward a more complete Usage Guide section. The Advanced Usage Guide and other sections will come at a later time.

Throughout the documentation, a parenthesized comment that starts with (.. indicates an editorial comment on the current status of implementation or documentation. In this version, however, an attempt has been made to remove all discussion of future capabilities and to document instead only the capability that is already there.

Source for code examples in the Usage Guide is included in the text where it occurs. Source consists of fragments taken from example programs. The example programs themselves may be downloaded from the web pages and compiled on your own system, using links provided for that purpose. (.. currently there are no code examples in the Usage Guide, which so far is only a general overview of the activity library ..)

There is also a directory of test programs (GridTurtle test programs
(../grid.html)
, contained within the documentation release directory) that provides additional code examples of many basic features of the defobj, collections, and activity libraries. These code examples are very rough and subject to change, however, since their main use is with each new release of the libraries to test various new or existing features. They are sometimes the only source, however, for examples of some library features that have not yet been used anywhere else.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Usage Guide

(.. This is an overview of activity library function. Code examples needed by a full usage guide will follow later.)

Role of the activity library in Swarm

The activity library provides a foundation for dynamic, object-oriented simulation in Swarm. Swarm assumes that a user defines an object-oriented representation for the structure of a world to be simulated. Once such a representation is established, activity library components are responsible for generating all state changes and flow of information within it. These state changes must be carefully controlled to ensure that they occur only at appropriate times and places in the model. The activity library provides mechanisms to establish this control.

Once the simulation of a dynamic model begins, everything that happens to the model occurs as a direct result of messages sent to world objects by activity library components. These components are also represented by objects, since this is the most effective means for their representation as well, but their purpose is not to represent the current state of the simulated world, but rather to generate changes in world objects by sending messages in a valid order. The activity library has two basic categories of components: those that represent messages to be sent, including constraints on permissible order for sending them, and those that execute the message sends these representations specify, making sure they conform with all constraints.

All changes to a model occur as a result of messages sent to objects within it. These messages invoke compiled methods of receiving objects, which may change local state or send messages to other objects to propagate effects as widely as needed. Only top-level function calls or message sends that initiate model processing are the concern of the activity library. Because the activity library finally interacts with the world model just by invoking its defined methods, the model can prepackage as much behavior as it can in the form of fast compiled methods. The activity library just triggers the prepackaged behavior at proper times under its own explicit, dynamically interpreted representation.

As conditions shift during execution of a model, the model can also examine and alter the representation of its future behavior contained in the activity structures that drive it. Any changes to this future behavior, however, still occur only as a result of some currently executing action initiated by the activity library.

Activity library components

The activity library defines one set of basic object types to provide a very rich representation of the kinds of message sending patterns it could generate. Its other major category of object types controls and tracks a current state of processing within these representations.

Both these representations are more abstract than typical object representations, since they deal not with any constant state which can be statically analyzed, but with shifting patterns of messages to be fired to generate such state. While the basic structure of a Swarm simulation is that of a discrete event simulation, its activity representation is also more complicated than most such systems.

There are two basic reasons for the potential complexity of a Swarm activity representation. One is that Swarm supports a decentralized representation of activity to be generated, both to reflect the nature of much of the behavior that it simulates, and so that execution may be distributed across multiple parallel processors. To avoid any need for synchronizing its decoupled activities more often than necessary, Swarm enables a model designer to avoid overspecifying constraints on the patterns of message sends which might potentially be valid. The partial order representation it adopts for a distributed and decoupled plan of activity is inherently more complicated than the single centralized event list adopted by many discrete event simulation systems.

The other reason for potential added complexity is that Swarm's representation of intended activity may be broken into many separate, modular components, which can be bound together in various ways to create larger components. The Swarm composition structure, when fully implemented, will be fundamentally as powerful as the modular abstractions found in most programming languages, with the added complexity of controlling the time at which various events occur.

In spite of this high potential complexity, none of these features is needed for many simple models, such as those that contain only a small variety of basic behaviors, or which define all their behavior to occur at regular, repeating timesteps. The Swarm representation is extremely rich to enable it to scale to large, multi-level models with a variety of dependent behaviors built into the model, and also to facilitate running models on massively parallel machines. Many of the features which seem complex, moreover, are also especially well-suited to building modular and reusable library components. It is anticipated that many of the more advanced features will find their heaviest usage in pre-built libraries that hide internal complexity from applications that use them.

No matter how complex the structures built from them, all the activity library components finally result in the direct execution of messages sends to objects in a model. Because of this direct execution, very precise meaning can be defined for each of them, in terms of message sends that can or must occur. Every component of a structure to be executed, and every event of its execution, can also be accessed using the interfaces of an object-oriented representation. This means that tools can be built to understand and interact with any components to be executed, and also to trace and debug activity as it occurs. None of these tools have been built yet as part of Swarm, but many of the activity library interfaces are present to support them, not because normal use of the library requires the added components.

Action plan components

The first set of activity library components represent messages to be sent to objects in a model, together with constraints in the order in which they may be sent. All these components are defined as subtypes of the one generic type called ActionPlan.

The two basic kinds of action plans are a simple group of actions to be performed in some order, called an ActionGroup, and a series of actions to be performed at discrete points in time, called a Schedule. The basic element of an action plan is a simple object called an Action. An action defines a particular message to be sent to an object or objects.

In its representation of action plans, Swarm relies heavily on the dynamic message sending capabilities of the Objective C language. The support of Objective C for dynamic message sends is absolutely crucial to Swarm's implementation of generic activity structures. Objective C defines a special kind of data value called a selector, which identifies a message according to its name. One of these selector values is stored in each action of an action plan. During execution of the plan, Objective C machinery performs the actual message send using its selector.

Action plans are free-standing objects that may be created directly by a user program, using a variety of selectable create-time options. Once created, action plans may also refer to each other, by containing special kinds of actions to start another action plan or to perform all the actions within it.

Individual actions can be created only as completely controlled components of some action plan. They are created not by a standard create message, but by special messages (each containing the phrase createAction in its name) sent to an action plan that creates the action as part of the action plan. If the same action is needed in more than one plan, it has to be created in each plan where it is needed. If the entire action plan is dropped, all its actions are also automatically dropped. So the only action plan components which need to be directly managed by an application are the action plans themselves.

Action plans are relatively straightforward components, because they are entirely passive. The only actions they ever contain are those which an application has created in them. These actions stay in them indefinitely unless a special option is requested to clean them out as each one is executed. Because these plans are passive, read-only components to all execution machinery, if the same pattern of actions needs to be triggered at multiple points in a model, it's perfectly valid to create one copy of the actions in an action plan, and then refer to the plan anywhere the actions may be needed.

Even though action plans are passive, containing only what has been placed in them, there is no requirement that their contents remain fixed. Both of the action groups and schedules are implemented directly as collections of their actions. New actions may be added at any time, and existing ones may be dropped or moved from position to another. The contents of action plans may be as dynamic as they need to be to represent the future actions needed in a model. None of these shifting contents has any effect on the model until actually processed during model execution.

Model execution components

The main responsibility of model execution is just to perform the actions specified by an action plan, in an order of execution consistent with any requirements of the plan. Each action plan may specify as few or as many constraints as it likes over the possible order in which its actions could be performed. Given these specifications, the execution objects are entirely responsible for selecting each action to be performed and then performing them.

There are two simple cases of ordering that account for most all usage, including that of the current Swarm demo programs. One is to permit the actions to be performed in any order, including concurrent execution if hardware and software support were available to do this. The other is to require them to be performed in some specific sequence, one after another, so that the effects of one action could depend on actions which preceded it. This sequence could be either fixed and predetermined, or dynamically established each time the action plan is performed. If a dynamically determined sequence is needed, perhaps even selecting which actions are to be performed at all, users can provide custom subclasses for an action plan and an execution object that performs it. A builtin option is to generate an entirely random sequence each time a plan is executed.

Each time an action plan is being processed, a special kind of object called an activity is created entirely automatically by the runtime processing machinery. These activity objects implement the internal machinery of a virtual processor which has the ability to execute action plans. To get any activity started on a model, a processor is first initialized to run a single top-level plan. All other activity during the lifetime of a model must occur as a result of actions initiated by this plan (which may include the startup of other plans).

Because the activity objects are internal to the processing machinery, an application can usually just ignore their existence. They come and go dynamically as various action plans are started and completed, all arranged in a stack or tree of current activities controlled by a single top-level processor. The activity objects are potentially useful, however, to obtain information about the context in which an action is currently running, or to build debugging or tracing tools to understand actions being performed.

One of the kinds of context information available from an activity object is the current time of a clock value incremented as a schedule is processed. A schedule is a kind of action plan in which all actions occur at specific points in time explicitly established within the schedule. As a schedule is executed, the activity object keeps a current time clock, which holds a global time from the start of all model execution, regardless of the time when the schedule itself was started. An activity object that processes a schedule is called a timeline activity, because its time continually increases from a global base time regardless of times contained in the schedule.

New activities are created whenever an action plan being processed contains an action to perform another action plan, or to start another action plan for autonomous execution. If one action plan performs another, its own processing is stopped until a new activity processing the other action plan completes. If one action plan starts another, the new action plan is started as an autonomous activity controlled only by a higher-level, containing activity.

A Swarm is an activity that exists only to control and coordinate other started subactivities. Unless the subactivities have some special form of explicit synchronization, none of their internal actions has any required ordering relative to those of other plans except as explicitly established during activity execution. The swarm can serve as a simple container of started subactivities which only occasionally synchronize for messages they send to each other. The swarm can be used to hold collections of objects as well as its subactivities as needed to help them coordinate with one another.

One special form of synchronization within a swarm is built into the virtual processing machinery. This synchronization interleaves the actions that occur for every successive time value during processing of timeline subactivities. This merging of actions is often relied on to interleave display and analysis processing with the scheduled actions of a base model. Since no other mechanisms for subactivity coordination are implemented in the current version of Swarm, synchronizing subschedule activities is the major current role of the Swarm activity type. In later versions of Swarm, a swarm will also serve as an important means for organizing a large models into clusters of more densely interacting components, and will also provide a basis of decomposition for parallel execution.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Advanced Usage Guide

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Interface Reference

Dependencies

Following are the other header files imported by <activity.h>:

```
#import <collections.h>
```

The activity library follows the Swarm library interface conventions
(../defobj/library.html)
of the defobj library
(../defobj.html)
. The activity library relies heavily on the basic collection types defined in the collections library
(../collections.html)
, and the collection interfaces are an integral part of the interfaces defined within this library. Initialization of the activity library is done automatically by the containing Swarm libraries, and automatically initializes the collections and defobj libraries as well.

List of Reference Pages

- ProcessType
(ProcessType.html)
- ActionType
(ActionType.html)
- CompoundAction
(CompoundAction.html)
- ActionGroup
(ActionGroup.html)
- Schedule
(Schedule.html)
- SwarmProcess
(SwarmProcess.html)

- Action Types
(Action.html)
- Activity Types
(Activity.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ProcessType – specification of a process

Synopsis

```
//
// ProcessType -- specification of a process
//
@deftype ProcessType
@end
```

Description

```
(..  For now, ProcessType is just a placeholder for future capability.  ..)
```

```
ProcessType will eventually define support for parameterization of all processes.
A process is a uniquely identified course of events that conforms to a specification
of external and internal behavior.
```

```
Processes include both objects and executable actions.  Objects change their state
and behavior based on a sequence of externally initiated actions.  An action, in contrast,
typically has no external behavior other than to be executed in its entirety as a
unit (yielding either an observable result or some change of state in a surrounding
environment).
```

```
Parameterization will be based a uniform framework that defines all input parameters,
internal state variables, and/or final results, and that provides for binding of these
data in a context of execution.  This parameterization framework is still being established,
but ProcessType objects will eventually record the specification of any process, to
the extent that this specification is defined.
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ActionType – specification of an executable process

Synopsis

```
//
// ActionType -- specification of an executable process
//
@deftype ActionType <ProcessType>- activateIn: swarmContext;
@end
```

Description

An action type is a type of process that may be initiated as a unit of execution by an external request. A typical action has a well-defined duration determined by a fixed set of actions that execute within it. Externally initiated interaction typically occurs only at the start or end of the overall process. A typical action is executed in its entirety once an external request that initiates it has occurred. Some actions may also have internal events that cannot begin or complete until other actions from a containing environment have also begun or completed their execution. Such ordering constraints can be defined either within an action type or as part of a dynamic context of execution.

Executable actions include both actions compiled in a host language (such as C functions or Objective C messages) and compound actions built at runtime for interpretation by the Swarm abstract machine.

(.. For now, the only subtype of ActionType is CompoundAction. Types for compiled actions such as functions and messages have not been defined yet. ..)

Inherited behavior

```
@deftype ActionType <ProcessType>
```

ActionType inherits standard parameterization support from ProcessType. (.. Currently, no such support is defined. ..)

ActionType behavior

`- activateIn: swarmContext;`

The activateIn: message is used to initialize a process for executing the actions of an ActionType. This process is controlled by an object called an Activity. The activateIn message initializes an activity to run under the execution context passed as the swarmContext argument, and return the activity object just created. If the execution context is nil, an activity is returned that allows complete execution control by the caller. Otherwise, the execution context must be either an instance of SwarmProcess or SwarmActivity. (These objects are always maintained in one-to-one association with each other, either one of the pair is equivalent to the other as a swarmContext argument.)

If a top-level activity is created (swarmContext is nil), the created activity may be processed using activity processing commands such as run, step, etc. If an activity is created to run under a swarm context, the swarm itself has responsibility for advancing the subactivity according to its requirements for synchronization and control among all its activities. Activating a plan for execution under a swarm turns over control to the swarm to execute the subactivity as a more-or-less autonomous activity.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

CompoundAction – collection of actions to be performed in some order

---

Synopsis

```
//
// CompoundAction --
//   a collection of actions to be performed in any order consistent with a
//   set of ordering constraints
//
@deftype CompoundAction <;ActionType, Collection>CREATING
- (void) setDefaultOrder: aSymbol;
- (void) setAutoDrop: (BOOL)autoDrop;
USING
- getDefaultOrder;
- (BOOL) getAutoDrop;
@end
```

Description

An compound action is the supertype of ActionGroup and Schedule. A compound action defines an executable process that is composed from the execution of a set of actions in some defined order.

Inherited behavior

```
@deftype CompoundAction <ActionType>
```

CompoundAction is not directly creatable. One of its subtypes must be created instead. ActionPlan inherits the basic ability to be activated for execution from ActionType
(ActionType.html)
.

Create-time options

```
CREATING
- (void) setDefaultOrder: aSymbol;
- (void) setAutoDrop: (BOOL)autoDrop;
USING
- getDefaultOrder;
- (BOOL) getAutoDrop;
```

DefaultOrder

The DefaultOrder option indicates the ordering to be assumed among actions of the plan when no other explicit ordering has been assigned. Beyond this initial ordering, additional ordering constraints can be added selectively using partial order specifications on individual actions. (.. Partial order order constraints are not yet implemented.)

The value for DefaultOrder is a symbol that may have one of the following values:

```
extern id <Symbol>  Concurrent, Sequential, Randomized;
```

The Concurrent value of the DefaultOrder option indicates that can actions be run in any order (including at the same time, if hardware and software to do this is available) without no impact on the net outcome of the actions. The claim that action results are independent of their execution order gives the processing machinery explicit leeway to execute them in any order it chooses. In the current implementation on a single, serial processor, actions are always processed sequentially even if marked concurrent, because that is the only way they can be. In future versions, however, special runtime processing modes may be defined even for a serial processor, which would mix up execution order just to confirm the independence of model results.

The Sequential value for the DefaultOrder option is the default. It specifies that the actions must always be executed in the same order as they occur in the plan. This order is ordinarily the same order in which actions are first created in the plan, unless actions are explicitly added elsewhere the collection that underlies a plan. This option is always the safest to assure predictability of results, but it excludes the ability to run the actions in parallel. To better understand and document a model design, it is worth annotating action plans with an explicit indication as to whether they do or do not depend on a Sequential order.

The Randomized value for the DefaultOrder option specifies that the model results do depend on execution order, but that the order in which the actions were created or added has no special significance. Instead, the method of dealing with order dependence is to generate a random order each time a collection of same-time actions is processed. The random order will be generated from an random number generator internal to the processing machinery.

AutoDrop

The AutoDrop option specifies that as soon as any action been processed by a running activity, the action is removed from the plan and dropped so that it will never appear again. This option is useful for plans that are created for a one-time use, never to be used again. This option is especially useful for a dynamic schedule that continually receives new actions to be executed at future times, but will never repeat actions that were previously scheduled. Depending on the underlying implementation of the schedule, making sure the old actions get dropped using AutoDrop can considerably improve the performance of the schedule.

When an option like AutoDrop is used, or whenever the contents of an action plan undergo a significant amount of dynamic update, the action plan would ordinarily be intended only for a single point of use. If AutoDrop is specified, a restriction against multiple active references is enforced. An error will be raised if two activities ever attempt to process a plan with AutoDrop enabled at the same time.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ActionGroup – collection of actions under partial order constraints

Synopsis

```
//
// ActionCreating -- protocol shared by ActionGroup and Schedule
//
@deftype ActionCreating
- createAction: anActionType;

- createActionCall: (func_t)fptr;
- createActionCall: (func_t)fptr : arg1;
- createActionCall: (func_t)fptr : arg1 : arg2;
- createActionCall: (func_t)fptr : arg1 : arg2 : arg3;

- createActionTo: target message: (SEL)aSel;
- createActionTo: target message: (SEL)aSel : arg1;
- createActionTo: target message: (SEL)aSel : arg1 : arg2;
- createActionTo: target message: (SEL)aSel : arg1 : arg2 : arg3;

- createActionForEach: target message: (SEL)aSel;
- createActionForEach: target message: (SEL)aSel : arg1;
- createActionForEach: target message: (SEL)aSel : arg1 : arg2;
- createActionForEach: target message: (SEL)aSel : arg1 : arg2 : arg3;
@end

//
// ActionGroup --
//   a collection of actions under total or partial order constraints
//
@deftype ActionGroup <CompoundAction, ActionCreating, OrderedSet, CREAT-
ABLE>@end
```

Description

An action group is an action plan whose basic representation is a sequence of actions that have been created within it. Even though a sequence is used to represent the actions of the collection, this does not necessarily mean that the execution machinery will process them in this order, since the permissible order of execution is still controlled by the DefaultOrder option inherited from ActionPlan. This initial order may also be further constrained by partial order constraints among the actions. (.. Partial order constraints are currently unimplemented.)

Inherited behavior

```
@deftype ActionGroup <CompoundAction, ActionCreating, OrderedSet, CREATABLE>
```

ActionCreating defines the createAction messages for ActionGroup just so that this protocol may be shared with Schedule, where they provide a convenience interface for the creation of actions in the schedule at time zero. These messages will be documented here in ActionGroup as if they had been defined on ActionGroup itself.

An action group inherits its underlying representation from the OrderedSet (../collections/OrderedSet.html) type of the collections (../collections.html) library. All the members of the ordered set must consist only of actions that are created by one of the createAction messages defined on ActionGroup itself. Once the actions are created, they may be accessed or traversed using standard messages of the OrderedSet type. Partial orders on these actions are also established by inherited messages. (.. Currently, partial orders are not implemented on either collections or action plans, but the messages inherited from OrderedSet will be the basic method for maintaining them once implemented.)

ActionGroup behavior

Actions are created within the ActionGroup using one of the createAction messages. A createActionTo: message specifies that the action to be performed is defined by an Objective C message selector to be performed on a receiving object plus any required arguments. The message selector is specified by the message: argument and the receiving object is specified as the first argument, target.

Given a variable aTurtle already initialized to a valid object, following is an example of a createActionTo message:

```
[createActionTo: aTurtle message: M(print)];
```

M() is a macro defined in defobj.h which is just a convenient shorthand for the Objective C expression @selector(message-name). As in any selector, the message name must include all parts of the message including colons and subsequent argument portions. An example of a message with one additional argument is the following:

```
[createActionTo: aTurtle message: $m(move:) : (id)10];
```

The createAction messages declare all arguments of the message to be of object id type, but you are free to cast other pointers and values up to the limits defined by the global portability assumptions (../defobj/globalport.html). These is not portable across all machine architectures, but is expected to be portable across the 32-bit and 64-bit architectures on which Swarm will be supported. The message you send must still be declared to receive the type of argument you actually pass, before you cast it to the id type.

(.. Alternative approaches to argument typing are currently in development, but these will supplement rather than replace the current forms of createAction messages.)

The createAction: message specifoes that processing of another action type is to be performed by the action. The referenced action type is performed in its entirety, from start to finish, as the effect of the single created action.

The createActionCall: messages are similar to the createActionTo messages, except they specify the action to be performed as a binding of a C function to a list of argument values. The correct number of arguments for the function pointer passed as the initial argument must be supplied.

The createActionForEach: messages define a message to be sent in the same way as the createActionTo messages, but they assume that the object passed as the target argument is a collection object. They specify that each object available from that collection, using the standard messages of the Collection
(../collections/Collection.html)
type in the collections library, is to receive the specified message.

The createAction: messages specify that processing of another action plan is to be performed by the action. The referenced action plan is performed in its entirety, from start to finish, as the effect of the single created action. If parameters are defined on the action plan, arguments for the action can also be given. (.. Currently this is not supported, so only the single argument form of createAction: should be used.)

Each of the createAction messages returns the action object which it creates. Each different kind of createAction message returns a different type of Action object with a matching name. These Action objects provide access to all the information with which the Action was initialized. The complete set of Action object types is defined below together with the messages that may be used to access their contents. (.. The implementation of the Action objects is currently undergoing change as the responsibility for parameter and return value typing gets taken over by ActionType in defobj.)

The action objects are an integral, controlled component of the action plan in which they are created. If they are removed from the action plan collection using a remove message, the only collection in which they may be reinserted is the same collection from which they came. It is permissible, however, to modify the base representation sequence by removing from one position and reinserting at another.

The contents of the action objects would ordinarily be referenced only by general-purpose tools to display or analyze an action plan structure. The action object id, however, can also be used to establish additional ordering constraints on any of the actions in a plan. These constraints supplement any initial order constraints established by the DefaultOrder option. Constraints may be established not only within the same plan, but between actions in different plans. (.. Partial order constraints are not yet implemented.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Schedule – collection of actions ordered by time values

Synopsis

```
//
// Schedule -- collection of actions ordered by time values
//
@deftype Schedule <CompoundAction, ActionCreating, Map, CREATABLE>CREATING
- create: aZone setRelativeTime: (BOOL)relativeTime;
- create: aZone setRepeatInterval: (timeval_t)repeatInterval;
SETTING
- (void) setConcurrentGroupType: groupType;
- (void) setSingletonGroups: (BOOL)singletonGroups;

- (void) setRelativeTime: (BOOL)relativeTime;
- (void) setRepeatInterval: (timeval_t)repeatInterval;
USING
- getConcurrentGroupType;
- (BOOL) getSingletonGroups;

- (BOOL) getRelativeTime;
- (timeval_t) getRepeatInterval;

- at: (timeval_t)tVal createAction: anActionType;

- at: (timeval_t)tVal createActionCall: (func_t)fptr;
- at: (timeval_t)tVal createActionCall: (func_t)fptr:arg1;
- at: (timeval_t)tVal createActionCall: (func_t)fptr:arg1:arg2;
- at: (timeval_t)tVal createActionCall: (func_t)fptr:arg1:arg2:arg3;

- at: (timeval_t)tVal createActionTo: target message: (SEL)aSel;
- at: (timeval_t)tVal createActionTo: target message: (SEL)aSel:arg1;
- at: (timeval_t)tVal createActionTo: target message: (SEL)aSel:arg1:arg2;
-
at: (timeval_t)tVal createActionTo: target message: (SEL)aSel:arg1:arg2:arg3;

- at: (timeval_t)tVal createActionForEach: target message: (SEL)aSel;
- at: (timeval_t)tVal createActionForEach: target message: (SEL)aSel:arg1;
-
at: (timeval_t)tVal createActionForEach: target message: (SEL)aSel:arg1:arg2;
- at: (timeval_t)tVal createActionForEach: target message: (SEL)aSel:
                                                                arg1:arg2:arg3;
```

```
@end
```

Description

A schedule is compound action whose basic representation is a sorted Map of actions that have been created within it. The key value associated with each of these actions is an unsigned integer value for which the typedef timeval_t

is supplied.

Inherited behavior

```
@deftype Schedule <CompoundAction, ActionCreating, Map, CREATABLE>
```

A schedule inherits its underlying representation from the Map
(../collections/Map.html)
type of the collections
(../collections.html)
library. All the members of the ordered set must consist only of actions that are created by one of the createAction messages defined on Schedule itself. Once the actions are created, they may be accessed or traversed using standard messages of the Map type. The key values of this collection, however, must be cast to and from the id type defined for key values by the Map type. Partial orders on these actions can also established by messages inherited from Map. (.. Currently, partial orders are not implemented on either collections or action plans, but the inherited messages will be the basic method for maintaining them once implemented.)

Create-time options

RelativeTime

The RelativeTime option specifies that all the times in the schedule are relative to the time when processing of the entire schedule begins. Otherwise, the times are assumed to be absolute times, with their base in the starting time of the entire model.

(.. This option is currently supported only on schedules, not swarms.)

RepeatInterval

The RepeatInterval option specifies that as soon as all actions in the schedule have completed, it is to be rescheduled at a time computed as the time at which the schedule was last started, plus the value specified as the RepeatInterval argument. This option overrides the normal default that times are considered absolute, as if the RelativeTime option had also been specified at the same time. All scheduled times must be less than the specified repeat interval, or an error will be raised. The RepeatInterval option can continue to be reassigned to different values after a schedule has been created, but the interval value must always be greater than the scheduled times of any actions which it contains.

(.. This option is currently supported only on schedule, not swarms.)

ConcurrentGroupType

The ConcurrentGroupType option is used to control handling of multiple actions which end up being scheduled at the same time value. As far the schedule representation is concerned, these actions are assumed by default to be concurrently executable, and the processing machinery is free to process them as such if no ConcurrentGroupType option is specified.

If a different interpretation of actions at the same time step is needed, the ConcurrentGroupType option may be specified. The argument of this option must be an object that when given a standard create: message will return an object having all the structure of a standard ActionGroup object. For example, if you want to specify that actions at the same time step are to be processed in random order, you could use the following procedure:

```
concurrentGroupType = [ActionGroup customizeBegin: aZone];
  [concurrentGroupType setDefaultOrder: Random];
  concurrentGroupType = [concurrentGroupType customizeEnd];

  aSchedule = [Schedule createBegin: aZone];
  [aSchedule setConcurrentGroupType: concurrentGroupType];
  aSchedule = [aSchedule createEnd];
```

Note that instead of using the normal createBegin/End sequence to create an instance, the concurrentGroupType is initialized using the customizeBegin/End sequence to customize a type without creating an instance. For more information, see the defobj library documentation.

In addition to overriding the standard ActionGroup type, the concurrent group type may be implemented by a custom subclass of the ActionGroup implementation which you supply yourself. A custom subclass is free to implement custom rules for how to combine all the actions which happen to arrive at the same time value. For example, it could decide that some actions are not to be executed at all, or it could create entirely new actions to be executed instead of or in addition to those which were originally scheduled. (.. Specific rules for writing a custom ActionGroup subclass will be documented at a future time, but all the apparatus to do so is present today. A concurrent action group is currently used to maintain the proper order of execution among subswarms of an owner swarm.)

SingletonGroups

The SingletonGroups option indicates that an action group should be created for every time value which is present, even when only a single action is present at the time value.

Ordinarily, a concurrent action group is created to process actions at the same timestep only if more than one action is scheduled at that timestep. The overhead of these action groups is relatively low, because it just creates a single new object to which actions are directly linked, but it is still faster to avoid creating them if only one action is present at a timestep. If a custom subclass is being provided however, it may need to examine the actions at a timestep even if there is only one.

---

timeval˙t – type for clock values

The following typedef is provided for convenience in declaring time values:

```
typedef unsigned   timeval_t;        // type for time values
```

Values of this type are used as keys when inserting actions into a schedule at a particular time, or for querying the current time value of a swarm or schedule during its execution. (The shorter name time_t has already been taken by one the standard C libraries.)

For very long running models using finely divided units of time, it is quite possible that a 32-bit unsigned time value could overflow. Special messages are provided in the execution machinery to reset all times to a different base value if a danger of overflow exists. If this machinery is exercised, all times of all referenced action plans are reset to a different base value in unison, which makes the shift as transparent as possible. Separate support is also available to declare time values that are subunits of the discrete clock values in a containing schedule or swarm. (.. Currently, all this extended time unit support is unsupported, though there is an example of scheduling at subunit times in a GridTurtle
(../grid.html)
test program.)

Schedule behavior

The messages to create actions within a schedule are essentially the same as those for ActionGroup

(ActionGroup.html)
, except for the presence of an initial at: argument indicating the time at which an action is to be performed. Except for the time associated with each action, meaning of the createAction messages is the same as for ActionGroup.

When multiple actions are all scheduled at the same time, they are all inserted into a concurrent action group created for that time value. The ConcurrentGroupType option may be used to override the default action group for these concurrent actions by a custom user-defined subclass. (.. Details of doing this are not yet documented, but there are examples.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

SwarmProcess – object that encapsulates a collection of processes

Synopsis

```
//
// SwarmProcess --
an object that holds a collection of concurrent subprocesses
//
@deftype SwarmProcess <ActionType, Zone, CREATABLE>CREATING
- setInternalZoneType: aZoneType;
- (void) setSynchronizationType: aScheduleType;

- (void) setInternalTimeMultiplier: (timeval_t)multiplier;
USING
- getSynchronizationType;
- (timeval_t) getInternalTimeMultiplier;

- getInternalZone;
- getActivity;
@end

//
// ActivationOrder -- default concurrent group type for swarm
//
@deftype ActivationOrder <ActionGroup, CREATABLE>@end
```

Description

A swarm is an object created to coordinate the processing performed on other action types. Just like any other action type, a swarm may be activated for running as a subactivity of another swarm. A swarm contains no actions of its own, since its only purpose is to coordinate the processing of other activities.

SwarmProcess is the type of an object that can be activated to run in some context, and which itself then defines a context in which other action plans can be activated. Subtypes of this generic type can each define a custom interface to activities running within the swarm. The generic swarm type defines only access to an underlying activity that keeps track of the activated or running subactivities, along with an option that may be set to a collection of objects currently being processed by the swarm.

If subactivities running within a swarm are indexed by time (running either schedules or other swarms), the default coordination performed by the swarm is to merge all the different scheduled actions so that the actions of one time value are all performed before the actions of any later time value. If the actions are not indexed by time, or if an subactivity has been released from strict time synchronization with other subactivities in the swarm, the default coordination is limited to partial order constraints among the action plans being run. (.. Currently only the merging of schedules and subswarms by time values is supported.)

A public superclass called Swarm is supplied in the objectbase library (../objectbase.html) for use in implementing custom swarm classes. Such a class can override the method for the activateIn: message of ActionType (ActionType.html) to automatically start subactivities to run within a swarm when the swarm is first activated.

Inherited behavior

```
@deftype SwarmProcess <ActionType, Zone, CREATABLE>
```

SwarmProcess inherits the messages of both ActionType (ActionType.html) and Zone (../defobj/Zone.html). Inheritance of zone behavior means that a swarm can be used as the argument of a create: or createBegin: message, for creation of an object within the internal zone of a swarm.

The default value for the SynchronizationType option is not a generic action group, but a special predefined subtype of ActionGroup called ActivationOrder. This concurrent group type guarantees that actions scheduled to occur at the same time from different action plans running in the same swarm are executed in the same order in which the action plans were first activated.

Create-time options

```
CREATING
- setInternalZoneType: aZoneType;
- (void) setSynchronizationType: aScheduleType;
- (void) setInternalTimeMultiplier: (timeval_t)multiplier;
USING
- getSynchronizationType;
- (timeval_t) getInternalTimeMultiplier;
```

InternalZoneType

The InternalZoneType option sets the type of zone which is created by the swarm to hold objects created within the swarm. If set to nil, no internal zone is created within the swarm, and all use of the swarm as if it were a zone raises an error. The default of this option is standard Zone type. (.. Since there is no other Zone type yet, there's no reason to set this option yet except to turn off the internal zone. ..)

SynchronizationType

Synchronization type sets the type of schedule which is used internally by the swarm to synchronize subschedules. Its default is a schedule with a concurrent group of ActivationOrder.

InternalTimeMultiplier

(.. This option will be used in the future to support an internal time unit within the swarm that can be more fine-grained than the context in which the swarm is included. ..)

SwarmProcess behavior

```
- getInternalZone;
- getActivity;
```

getInternalZone returns a Zone object that is used by the swarm to hold its internal objects. Even though the swarm itself inherits from Zone and can be used as a Zone for nearly all purposes, this message is also provided so that the zone itself can be obtained independent of all zone behavior.

getActivity returns the activity which is currently running of subactivities within the swarm. This activity is the same as the value returned by activateIn: when the swarm was first activated. It returns nil if the swarm has not yet been activated.

Unlike other action types, swarms and swarm activities always exist in a one-to-one relationship, provided that the swarm has been activated. This restriction to a single activity enables the swarm to do double-duty as a custom object that provides its own interface to the activities running within the swarm.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

Action – element of a compound action

Synopsis

```
//
// Action -- messages shared by ActionCall, ActionTo, and ActionForEach
//
@deftype Action <GetOwner>- (int) getNArgs;
- (void) setArg1: arg1;
- getArg1;
- (void) setArg2: arg2;
- getArg2;
- (void) setArg3: arg3;
- getArg3;
@end

//
// ActionCall -- an action defined by calling a C function
//
@deftype ActionCall <Action>- (void) setFunctionPointer: (func_t)fptr;
- (func_t) getFunctionPointer;
@end

//
// ActionTo -- an action defined by sending an Objective C message
//
@deftype ActionTo <Action>- (void) setTarget: target;
- getTarget;
- (void) setMessageSelector: (SEL)aSel;
- (SEL) getMessageSelector;
@end

//
// ActionForEach --
an action defined by sending to every member of a collection
//
@deftype ActionForEach <ActionTo>@end
```

Description

Action is a common supertype of all action types which may be created within an action plan. Each action is always controlled by a single action plan to which it belongs. This action plan is referred to as its owner. Given the action object, its owner plan may be obtained using the inherited getOwner message.

Actions are allocated in the same zone as their owner plan, and may be created only using one of the createAction messages on an ActionGroup or Schedule. Each of these messages returns the action that was created as its return value. Actions in an action plan may also be obtained by processing the plan using any of its messages inherited from its underlying collection. Actions may also be removed from an action plan using a remove message on the underlying collection.

(.. Note: currently, an action cannot be removed while it is currently being executed. This means that the function or message called by an action cannot itself remove that same action from its action plan. This restriction will be removed in the future.) No matter what the type of action, the default ordering constraints established by its membership in an owning plan may be supplemented by explicit specification. These specifications are made directly between actions in the same plan, or between actions in different plans if they are suitably qualified. Given any action belonging to any plan, the set of specific explicit ordering constraints in which it participates may be obtained by one of the messages getSuccessors, getPredecessors, getCostarts, and Coends inherited from the KeyedCollection
(../collections/KeyedCltn.html)
type.

Separate subtypes of Action are defined for each of the various forms of createAction messages that create them. The current representation of these actions will be undergoing change as support for their parameter and return types is migrated into more basic support from the defobj
(../defobj.html)
library. Each action type provides messages to retrieve and set the values of all argument values bound into the action. These capabilities will remain, but different messages will eventually be supported. This documentation will be completed once the messages supported on Action types are finalized.

Roger Burkhart <rmb@santafe.edu>
(mailto:rmb@santafe.edu)

Last modified: 1997-12-19 (mgd)

Activity Types

Contents

The following object types are documented below:

- Activity
- GroupActivity
- ForEachActivity
- ScheduleActivity
- SwarmActivity

Name

Activity – state of processing within an action plan

Synopsis

```
//
// Activity -- state of processing within an action plan
//
@deftype Activity <DefinedObject, Drop>- run;
-               stop;
- (void) terminate;

- next;
- step;
- stepEntry;
- stepExit;

- getStatus;
- getHoldType;

- getAction;
- getActionType;
```

```
- (void) setOwnerActivity: ownerActivity;
- getOwnerActivity;

- getSubactivities;

- getControllingActivity;
- getTopLevelActivity;
- getSwarmActivity;
- getScheduleActivity;

- setSerialMode: (BOOL)serialMode;
- (BOOL) getSerialMode;

- getCurrentSubactivity;  // serial mode only
@end
```

Description

A object of type Activity implements the processing of actions within an action plan. Each subtype of action plan has a corresponding subtype of Activity that implements processing of its respective type of plan.

The Activity types are part of the virtual processing machinery of an action plan interpreter. All the important elements of this machinery, including their current state of processing, are exposed to aid in development of general-purpose tools such as debuggers. Except for applications that need to create and run their own reflective activities, direct use of activity types by a modeling application should be rare.

An new activity object is created by each activateIn: message sent to an action type. activateIn: initializes a new activity object and prepares it for processing, but does not itself execute any actions. Subsequent messages, such as run, must be used to initiate processing within an activity. If an activity is activated to run under a swarm context, the owning swarm itself controls all processing within the subactivity.

An Activity type holds a tree of currently running, or a potentially runnable, subactivities. Each activity object represents a state of processing within a single action plan. The structure is a tree because multiple activities could be running or eligible to run at the same time. This occurs whenever two activities are specified as concurrent, either because are being performed by concurrent actions in some other plan, or because they were started for autonomous execution under the same swarm. When parallel processing machinery is present, each activity could also be advancing its own state independent of those of any other activity.

The current implementation supports only a single serial processor. Under serial execution, only one activity may be active at one time. This means that the current state of processing may be represented by a single stack of current activities that traces a path to a single leaf of the runnable activity tree. When running in serial mode, messages are available to obtain the currently running leaf activity or useful context information such as the current time available from it.

Activity behavior

(.. Specific messages for activity objects are still subject to change as some of their underlying support migrates to the defobj

(../defobj.html)
library. For now, this section will document just the basic messages to start and control processing within an activity.)

```
- run;
-               stop;
- (void) terminate;
```

The run message continue processing of an activity from any state in which its execution has been suspended. An error is raised if the activity is already running. run returns either a Stopped or Completed status that the activity has when it is no longer eligible to be run further.

The stop message causes the a currently running tree of activities to stop further processing and return a Stopped status. terminate also stops a running tree of activities, but sets all activities within the tree to a status of Completed whenever the tree is next run. terminate may be used on either a running or stopped tree of activities. It is the standard way to terminate schedule that is endlessly repeating under the RepeatInterval option.

```
- next;
- step;
- stepEntry;
- stepExit;
```

The next executes actions within a single compound action while skipping over all processing of any complete action plans executed by those actions. The step message executes a single action within a tree of activities. The stepEntry message executes multiple actions within a single currently runnable activity, but stops the activity when a new level of action plan has just been started within the tree of activities. stepExit executes multiple actions in the same action plan until the activity for that action plan has completed, or cannot otherwise be run further. (.. Currently, stepEntry and stepExit are unimplemented.)

```
- getStatus;
- getHoldType;
```

The getStatus message returns one of the following codes for the current run status of a particular activity:

```
extern id <Symbol>  Initialized, Running, Stopped, Holding,
```

```
            Released, Terminated, Completed;
```

The getHoldType returns a code for the particular hold constraint under which an activity is currently holding. It returns nil if the basic status of the activity is not Holding. (.. Currently no hold constraints other than merging within a swarm are supported, and this message always returns nil.)

```
extern id <Symbol>  HoldStart, HoldEnd;
```

Activity context macros

The current running activity is always available for reference within any action running within the activity. (.. Currently, the getCurrentActivity() macro raises an error, but the other macros work as long as there is a currently running activity. ..) The following macros provide access to various information available from the current running activity and its owner activities:

```
#define  getCurrentActivity()
#define  getCurrentAction()
#define  getCurrentTime()
#define  getCurrentScheduleActivity()
#define  getCurrentSchedule()
#define  getCurrentSwarmActivity()
#define  getCurrentSwarm()
#define  getCurrentSwarmActivity()
#define  getTopLevelActivity()
```

Activity subtypes

(.. Only a synopsis of Activity subtypes is provided here for now. Note, however that access to a current time is available from a schedule or swarm activity, along with the ability to run a schedule until a particular time value is first reached.)

Name

GroupActivity – state of processing within an ActionGroup

Synopsis

```
//
// GroupActivity -- state of processing within an ActionGroup
//
@deftype GroupActivity <Activity>@end
```

Name

ForEachActivity – state of processing within an ForEach action

Synopsis

```
//
// ForEachActivity -- state of execution within a ForEach action
//
@deftype ForEachActivity <Activity>
- getCurrentMember;
@end
```

Name

ScheduleActivity – state of execution within a Schedule

Synopsis

```
//
// ScheduleActivity -- state of execution within a Schedule
//
@deftype ScheduleActivity <Activity>
- (timeval_t)   getCurrentTime;

- stepUntil: (timeval_t)tVal;  // unimplemented
@end
```

Name

SwarmActivity – a collection of started subactivities

Synopsis

```
//
// SwarmActivity -- a collection of started subactivities
//
@deftype SwarmActivity <ScheduleActivity>
@end
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Subclassing Reference

Subclassing is supported by the activity library as an integral technique for extending the framework it implements. There are two specific places where subclassing is expected as the normal technique of extension: definition of concurrent action groups, and definition of customized swarm objects. Concurrent action groups are not yet fully documented because their full interface is still being finalized. Swarm subclasses should inherit from the Swarm superclass defined in the objectbase (../objectbase.html) library; the activity library provides the underlying support packaged by this superclass.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Interface Design Notes

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Activity Library: Implementation Notes

(This section of documentation is not yet available.)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08 (mgd)

Objectbase Library

Documentation and Implementation Status

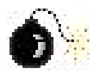

Compatibility

- 1.0.4 => 1.0.5
  No changes.
- 1.0.3 => 1.0.4
  The name of this library is now **objectbase**, it is has been renamed from **swarmobject** largely to reflect the more generic nature of the library and also motivated by the impending port of Swarm to Windows NT (to avoid filename conflicts with the `SwarmObject` class). There should be little, or no effect on the user, the only visible change is the fact that the actual library (.a) or (.so) file will now have a different name and the header file name has changed. A symbolic link from `objectbase.h` to `swarmobject.h` has been provided in the distribution, to ensure backwards compatibility, however, users should not continue to rely on this being so. Users should port references to `swarmobject.h` in their code to `objectbase.h`, because this symlink will be removed in a future release.
  Note: this is no way affects the `SwarmObject` class which remains the same as in all previous releases.
- 1.0.0 => 1.0.1
  The interface has changed again! *EmptyProbeMap* is now a subclass of *CustomProbeMap*, which is subclassed from *ProbeMap*. And a shortcut `create:` method was added to that branch.
  Also, a new method was added to *ProbeLibrary* called `isProbeMapDefinedFor` that serves to non-invasively test for the existence of a *ProbeMap* for a given class.
- *Beta* => 1.0.0
  The new interface for the *swarmobject* library might cause some problems for apps that worked under the Beta release of Swarm. To get the whole scoop, read the Library Interface Conventions
  (defobj/library.html)
  . To get a few hints at explicit changes you might need to make, check out: Porting apps to use the new *swarmobject* interface
  (objectbase/PortingAppsToNewInterface.html)
  .

Overview

The *objectbase* library encapsulates various fundamental aspects of the Swarm object model and defines the probing machinery used to take data from Swarm objects. Most of the underlying functionality of the classes defined here is contained in the *defobj*

(defobj.html)
and *activity*
(activity.html)
libraries.

The main classes defined here that are intended to be subclassed by users are Swarm and SwarmObject. The probing functionality provided here is mainly for use within the Swarm kernel. However, this probe machinery should be used when designing any interface between Swarm and any other agent or device.

SwarmObject

SwarmObject is the base class for all interactive agents in Swarm. It defines the standard behavior for Swarmstyle agents, which includes hooks for creation, probing, zoned memory allocation, and destruction.

Swarm

The *Swarm* class encapsulates all of the abstract execution machinery in the *activity* library, as well as the notion of a group of related objects. Metaphorically, a "Swarm" is a combination of a collection of objects and a schedule of activity over those objects. Hence, the *Swarm* class provides the behavior necessary for creating such an object and starting up the schedule.

Probes

Probes are idealized entities that are intended to allow the user to monitor and modify the innards of objects without explicitly providing the functionality to do so at compile time. Hence, they allow dynamic interaction with an objects instance variables and methods. Most of the functionality of probes is implemented here; but, at present, they are intimately linked to the *simtools*

(simtools.html)
library, which contains all the widgetry needed to use probes from a GUI. Despite this intimacy, probes are intended to be a general purpose mechanism for any agent or device to interact dynamically with Swarm objects. This Introduction to Probes

(probing.html)
will explain more about the reasoning and purpose behind probes.

Documentation

The *objectbase* library is defined and documented following the basic structure of the Library Interface Conventions

(defobj/library.html)
. The following standard sections of documentation are available:

- Interface Reference
(objectbase/Interface.html)
- Usage Guide
(objectbase/UsageGuide.html)
- Advanced Usage Notes
(objectbase/AdvancedUsage.html)
- Interface Design Notes
(objectbase/InterfaceDesign.html)
- Implementation Notes
(objectbase/Implementation.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Objectbase Library: Documentation and Implementation Status

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Swarmobject Library: Porting Apps to the New Interface

The problems that might arise because of this new interface should have to do with the placement or existence of C preprocessor `#import`s. In the beta release, the `objectbase.h` file simply included all the header files for each class defined in the objectbase library. These files contain the `@interface` declarations for each respective class. Hence, before, if an application file included `objectbase.h`, it was including the interface for every class in the library.

The Swarm standard module interface uses Objective-C protocols to establish which messages may be sent to objects of a class. Doing this isolates the classes in the library more completely from the rest of the code. So, the problems that are likely to arise from this new interface are ones of access to objects that reside in the objectbase library. And they are of two types:

1. Messages that are no longer available on one or more classes where they were available before. I.e. you see warnings like the following when you compile your app:

```
warning: 'ClassName' does not respond to 'messageName'
```

What this means is that since `objectbase.h` simply imported all the class header files, somewhere in that chaotic mess of an import tree, the header for the `ClassName` was imported. Now, since `objectbase.h` only defines protocols for the classes it contains, it no longer imported that header file.

Solution #1: If `ClassName` is a member of the objectbase library, then all you need to do is import `objectbase.h` in the file where that class (or an instance of it) is used. If, however, the class in question is a member of some other library, then you should find out which library it is a member of and include the appropriate header file. (This latter case is actually part of the justification for making these libraries into "modules" like this.... to avoid incestuous include trees.)

Another case where this occurs is when using objects that have been statically typed with a class name. In using the protocols instead of direct class interfaces, we are depriving the compiler of the ability to analyze the type of an object and decide if that object can accept a given message. This is a bit counter intuitive for those of us reared with strongly typed languages and styles. But, Swarm takes advantage of the freedom this allows and it is part of Swarm-style programming
(../defobj/style.html)
.

Solution #2: If this is what's happening to generate the above warning, there are two further degradations of typing available. The next step down is to take away the specific class type and replace it with a protocol qualification on a pure `id` type. For example:

```
ProbeMap * myProbeMap;
```

can be changed to

```
id <ProbeMap> myProbeMap;
```

This both allows the freedom from the static typing and adds to the readibility of the code. However, this still results in compile time checking for messages sent to that object with the protocol designated. So, any message you intend to send to that object must be defined or inherited by the given protocol. So, in some cases, one may want to go with the ultimate freedom in

```
id myProbeMap;
```

which, obviously, means the compiler will not second guess the programmer in what messages can be sent to that object. And if your variables are named appropriately, this doesn't really detract from the readibility of the code.

2. When subclassing from a class defined in the Swarm Kernel, it is necessary to import the `@interface` of that class. Before, when `objectbase.h` included all those interfaces, it was the only file needed if you wanted to subclass from a class in objectbase. Now, however, simply importing `objectbase.h` doesn't help at all. So, in order to subclass from a class that's a member of objectbase, you need to import the specific interface file to that class.
   The solution: If you had an app that worked before when you subclassed from something (say, Swarm) in objectbase but doesn't work now, simply replace the

```
#import <objectbase.h>
```

in the header file for the subclass with

```
#import <objectbase/Swarm.h>
```

Note: In order to reduce the number of broken apps due to this, we've explicitly imported `objectbase/SwarmObject.h` and `objectbase/Swarm.h`at the end of the `objectbase.h` file. This allows the subclassing of SwarmObject or Swarm by an import of `objectbase.h` without the explicit import of the interface files. But, this was done for backward compatibility and may go away in the future. So please subclass by explicit importation of the interface file of the class.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Objectbase Library: Interface Reference

The *objectbase* library uses the module definition conventions
(../defobj/module.html)
established by the *defobj* library
(../defobj.html)
. It also depends on standard supertypes and classes defined by this library.

Dependencies

Following are the other header files imported by <objectbase.h>:

```
#import <defobj.h>
#import <objectbase/SwarmObject.h>;
```

The *defobj* library interface is included to provide the basic object support. The *SwarmObject* class header file was provided for backward compatibility since the *objectbase* library has been brought into agreement with the library standards.

Object Types (with links to reference pages)

```
SwarmObject
 (SwarmObject.html)
                  /     |     \
                 /      |      \                   ProbeLibrary
 (ProbeLibrary.html)
                /       |       \
  ActivityControl
 (ActivityControl.html)
      Probe
 (Probe.html)
      ProbeMap
 (ProbeMap.html)
                      /   |        |  \   \
                    /     |        |   \    \
                  /       |        |    \     \
                /         |        |     \      \
        VarProbe
 (VarProbe.html)
       Message
 (MessageProbe.html)
   Custom
```

```
    (CustomProbeMap.html)
       \      Complete
    (CompleteVarMap.html)
                                Probe
    (MessageProbe.html)
       Probe
    (CustomProbeMap.html)
        \       Var
    (CompleteVarMap.html)
                                          Map
    (CustomProbeMap.html)
          \       Map
    (CompleteVarMap.html)
                                                  Complete

    (CompleteProbeMap.html)
                                                   Probe

    (CompleteProbeMap.html)
                                                    Map

    (CompleteProbeMap.html)
```
(don't want a graphic with image map to avoid excessive download times!)

Contents:

(../objectbase.html)

(../objectbase.html)

(UsageGuide.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

SwarmObject Class

```
//
// SwarmObject --
//   Superclass of most objects in a Swarm simulation.
//   Provides support for probing.
//
@protocol SwarmObject <Create, Drop>
USING
- (const char*)getInstanceName;
- getProbeMap;
- getCompleteProbeMap;
- getProbeForVariable: (const char *)aVariable;
- getProbeForMessage: (const char *)aMessage;
@end
```

A SwarmObject is an object that is intended to be a member of a Swarm. That is, its behavior will be perpetuated by messages sent from the schedule of events defined in the context of a Swarm object.

The SwarmObject is where the models of all the agents of a simulation will reside. Hence, most of the burden on defining the messages that can be sent to any agent lies with the user. SwarmObject inherits its basic functionality from the Create and Drop object types defined in the defobj library. It is useful to mention these messages here to show the complete span of messages that can dictate the behavior of an instance of a SwarmObject; but it is not useful to explain their function in detail. The messages SwarmObject inherits from defobj are:

- DefinedObject
  - - getType;
  - - getClass;
  - - (BOOL)respondsTo: (SEL)aSel;
  - - describe: anOutputStream;
  - - (BOOL)setDisplayNameTo: (const char *)aName;
  - - (const char *)getDisplayName;
  - - createIDString: aZone;
- Create
  - + create: aZone;

- + createBegin: aZone;
- - createEnd;
- + customizeBegin: aZone;
- - customizeEnd;
- - customizeCopy: aZone;

- Drop

  - - getZone;
  - - (void)drop;

In addition to these, SwarmObject is augmented by the messages that provide the probe support.

getInstanceName is a message used by the probe machinery to get a meaningful word or phrase to indicate the object being probed (and, in the case of a probe display, place it in a widget on the screen). Support for this has recently been added to the DefinedObject object type, thereby making it more generally available. So, this message will likely become obsolete soon.

getProbeMap returns a pointer to the ProbeMap for the object if there has been one created for that object's class. If it hasn't, then it create's a default one.

getCompleteProbeMap returns a newly created CompleteProbeMap for the object.

getProbeForVariable: (const char *)aVariable returns the VarProbe indexed in the ProbeMap by the string aVariable.

getProbeForMessage: (const char *)aMessage returns the MessageProbe indexed in the ProbeMap by the string aMessage.

(Interface.html)

(../objectbase.html)

(ProbeLibrary.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

ProbeLibrary Class

```
//
// ProbeLibrary --
// A (singleton) Class, whose instance is used as a container for aglobal
// mapping between classnames and their 'default' ProbeMaps. These defaults
// can be changed by the user, thus allowing him/her to customize the default
// contents of the ProbeDisplays generated when probing objects.
//
@protocol ProbeLibrary <Create, Drop, ProbeConfig>
USING
- setDisplayPrecision: (int) nSigDisplay;
- (int)getDisplayPrecision;
- setSavedPrecision: (int) nSigSaved;
- (int)getSavedPrecision;
- (BOOL)isProbeMapDefinedFor: (Class)aClass;
- getProbeMapFor: (Class)aClass;
- getCompleteProbeMapFor: (Class)aClass;
- getCompleteVarMapFor: (Class)aClass;
- getProbeForVariable: (const char *)aVar inClass: (Class)aClass;
- getProbeForMessage: (const char *)aMessage inClass: (Class)aClass;
- setProbeMap: aMap For: (Class)aClass;
@end

extern id <ProbeLibrary> probeLibrary;
```

The normal Swarm simulation will probably only ever contain one instance of this class, namely the *probeLibrary* object. This object is used for Library Generation of *Probe*s and *ProbeMaps*: its role is to cache one unique "official" *ProbeMap* for every Class ever probed during a run of Swarm. These *ProbeMap*s are generated lazily as they are requested.

- setDisplayPrecision: (int)nSigDisplay
  Sets the number of significant digits for floating point and double floating point numbers displayed on GUI widgets. Currently only implemented for `VarProbes` not `MessageProbes` as yet. This allows *all* probes checked out the global ProbeLibrary instance to access this displayed precision. However, individual probes can vary from this global default, by using the `setFloatFormat` method on a exisiting probe. See probing precision (../probing.html#precision) for more details on how to do this.

- (int)getDisplayPrecision

Gets the current display precision set in the ProbeLibrary instance. Mostly used internally.

- setSavedPrecision: (int)nSigSaved
  Sets the number of significant digits saved for floating-point and double floating-point numbers through ObjectSaver. This function sets the global default precision for all floating (and double floating) point numbers. This floating point precision affects all numbers saved via the ObjectSaver class. There is currently no way to override this global default for an individual probe.

-(int) getSavedPrecision
  Gets the current saved precision set in the ProbeLibrary instance. Mostly used internally.

Note: if neither of these methods are used to set the defaults, the precision in both cases defaults to six (6) significant figures, which should simulate the default precision before the introduction on this feature *on most platforms*. Previously, the sprintf "%g" default format was used to print out floats and doubles, which is not considered to have entirely portable behavoir.

- (BOOL) isProbeMapDefinedFor: (Class) aClass returns True
  If there is a non-nil value in the *ProbeLibrary* for that class and False otherwise.

- getProbeMapFor: (Class)aClass
  returns a *ProbeMap* for the *aClass* class. If a specific *ProbeMap* has been designed and installed in the *ProbeLibrary* for that class, then that *ProbeMap* is returned. If not, then a *CompleteProbeMap* is created and returned.

- getCompleteProbeMapFor: (Class)aClass
  This method returns a *ProbeMap* containing *Probe*s for all the instance variables and messages of the given Class (including inherited variables and messages). Since we are assuming that normal, single-level *ProbeMap*s will be more widely used, the current implementation of *ProbeLibrary* does not cache *CompleteProbeMap*s. However, if we find that there is a demand for *CompleteProbeMap*s, we will cache them as well.

- getCompleteVarMapFor: (Class)aClass
  This method returns a *ProbeMap* containing Probes for all the instance variables of the given Class (including inherited variables) but does not include any *MessageProbe*s.

- getProbeForVariable: (const char *)aVariable inClass: (Class)aClass

- getProbeForMessage: (const char *)aMessage inClass: (Class)aClass
  These methods are used to "check out" the appropriate Probes from the probe library. Note that the returned probes will be cached, so you must make sure to clone them if you want to make modifications to the probes without affecting the result of future requests for those same probes.

- setProbeMap: aMap For: (Class)aClass

In those cases where the user wants to set the standard probe map directly, she can provide the probe map using this method, and it will be cached as though it were produced by the library itself.

(SwarmObject.html)

(Interface.html)

(Probe.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

ActivityControl Class

```
//
// ActivityControl -- This class provides an object interface to
//                    an activity.  It contains a subset of the messages
//                    that any activity responds to and provides explicit
//                    control over that activity.
//
@protocol ActivityControl <SwarmObject>
CREATING
- attachToActivity: anActivity;
USING
- run;
- stop;
- next;
- step;
- stepUntil: (timeval_t)stopTime;
- (void)terminate;
- updateStateVar;
- getStatus;
- (const char *)getInstanceName;
@end
```

The ActivityControl class specifies an object that can be attached to an activity (regardless of how or where that activity is created) for the purpose of explicitly controlling the execution of actions on that activity's action plan. There is nothing that available through this class that is not already available through the variables or messages available via the activity itself. However, it packages up the main control messages and makes them available to other objects that may need control over an activity, thereby shielding the activity from directly receiving messages from outside objects and saving the user from having to parse the more complex interface to the activity.

Create Phase Protocols

attachToActivity: anActivity sets an instance variable inside the ActivityControl object that points to the Activity to be controlled. It then creates a Schedule upon which it places a message to itself to update its own variables.

Usage Phase Protocols

run sends a run message to the activity if the conditions are appropriate. This message causes the activity to continue executing the actions on it's schedule until either no other actions are waiting, at which point it returns Completed, or until it is stopped by a subactivity or a stop message to the activity.

stop sends a stop message to the activity if the conditions are appropriate. This message causes the control to move back up the run-stack and resume at the place in the code where the run was first executed, which doesn't mean much in a parallel execution context; but, in a serial context, it means that the next action on the super-activity will begin without finishing the rest of the current activity's actions.

next sends a next message to the activity if the conditions are appropriate. It runs an activity forward through as many actions as necessary until it hits a breakFunction, at which point it walks back up the tree of activities and returns Stopped. In most cases, this means that an entire action or action group on the activity under control will be executed, including completion of all subactivities.

step sends a step message to the activity if the conditions are appropriate. It causes the execution of a single action.

stepUntil: (timeval_t)stopTime sends a stepUntil: message to the activity if conditions are appropriate. This causes all actions on the activity's schedule, including any actions on subactivities' schedules, to be executed until the activity's relative time is equal to stopTime - 1.

(void)terminate sends the terminate message to the activity, which causes all actions and action groups to be removed from its schedule. This can cause some problems. If terminate is sent to an activity and, subsequently, a run is attempted on that activity, the program will exit with an error. A way to avoid this might be to eliminate the ActivityControl attached to the terminated activity; but, that would preclude the reuse of ActivityControllers on other activities. Ideally, a controller could be adjusted to monitor any of several activities and it could leap from one activity to another. So, this behaviour is advertised as a feature.

updateStateVar updates the ActivityControl instance variables and tests for the continued existence of the activity that's being controlled. This message is sent on each cycle of the schedule for the activity being controlled.

getStatus returns the status of the activity.

(const char *)getInstanceName This method is provided for backwards compatibility to the probe machinery. It echos the displayName resident in the object name database. When the probe machinery is upgraded to use this database, this method will become obsolete.

(CompleteVarMap.html)

(Interface.html)

(UsageGuide.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Probe Class

```
//
// Probe --
//   Abstract superclass of both VarProbe and MessageProbe.
//
@protocol Probe <SwarmObject>
CREATING
- setProbedClass;
SETTING
- setSafety;     // There should be a setSafety: BOOL
- unsetSafety;   // and getSafety methods instead...
USING
- clone: aZone;
- (Class)getProbedClass;
- (const char *)getProbedType;
@end
```

A Probe is simply an object that contains pointers to an element (instance variable or message description) of another object. The Probe contains instance variables that describe the referent's class and type. It's actually an abstract class that is further subdivided into VarProbe and MessageProbe, which represent the two basic types of elements of any object. The Probes are collected into a ProbeMap and subsequently installed in the ProbeLibrary.

Create Phase Protocol

setProbedClass: (Class)aClass sets the class of the referent and must be called at create time.

Setting Phase Protocol

setSafety turns on the option of checking the compatibility of the class of the object before any actions are performed on the referent. This is only used in the VarProbes and should probably be moved.

unsetSafety turns the safety option off.

Regular Protocol

(Class)getProbedClass returns the class of the referent as a Class pointer.

(const char *)getProbedType returns the typing of the probed variable/message. The typing is represented using the string-format provided by the Objective-C runtime system ('c' for char, '@' for id, and so on). In general the programmer will know the typing of the variable so this method will probably only be used internally by Swarm kernel code.

clone: aZone returns a clone of the probe. It is useful in case you got the probe by Library Generation or by the default version of Object Generation. In both cases you will have been given a unique shared instance of the probe - so if you wanted to make changes to the safety status of the probe etc. without affecting the other potential users of the probe, you would first clone it and then make the required changes.

(ProbeLibrary.html)

(SwarmObject.html)

(VarProbe.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

VarProbe Class

```
//
// VarProbe --
// Allows the user to inspect a given variable in any candidate
// which is an instance of, or inherits from, a given class
//
@protocol VarProbe <Probe>
CREATING
- setProbedVariable: (const char *)aVariable;
SETTING
- setStringReturnType: returnType;
USING
- (const char *)getProbedVariable;
- setFloatFormat: (const char *)format;
// Currently, setNonInteractive must be used *after* create phase
// In any case these should also become setInteractive: BOOL etc.

- setNonInteractive;
- (int)isInteractive;

- (void *)probeRaw: anObject;
- (void *)probeAsPointer: anObject;
- (int)probeAsInt: anObject;
- (double)probeAsDouble: anObject;
- (const char *)probeAsString: anObject Buffer: (char *)buffer;
- (const char *)probeAsString: (id) anObject Buffer: (char *)buf withFullPre-
cision: (int)precision;
- setData: anObject To: (void *) newValue; // pass by reference.
- (int)setData: anObject ToString: (const char *)s;
@end
```

This is a specialized subclass of the abstract class *Probe*. It completes the specification of a probe whose referent is an instance variable element of an object.

Create Phase Protocol

- setProbedVariable: (const char *)aVariable sets the variable being probed (elsewhere called the referent). The *aVariable* identifier is simply a character string consisting of the identifier of the variable referent. For example, if we were probing a variable called *unhappiness*, then the string would be "unhappiness". This method must be called during the create phase.

Set Phase Protocol

- setStringReturnType: returnType sets the format that will be used to print the variable.

Usage Protocol

- setFloatFormat: (const char *)
  Given an sprintf style formatting string (eg. "%.3f") to a floating-point (or double-floating-point) probe will use this formatting *only* when probing on a GUI display widget. See section on probing precision (../probing.html#precision) when probing floating point number in the documentation on probes.
- setSafety
- unsetSafety
  In order to ensure efficiency, the default behaviour of the *VarProbe* is to "trust" that the object it is being presented with is of an appropriate class (the one specified in *setProbedClass*, or a subclass of it). However, by calling *setSafety*, we force the probe to inspect the class of the target object before probing it.
- setNonInteractive
- (int)isInteractive
  By setting a *VarProbe* to be non-interactive we ensure that the user will not be able to change the value of the probe, but only observe it. The default is for variables of type (char *), char, unsigned char, int, unsigned int, floatand double to be interactive and for all the others to be non-interactive (with the exception of the id type this follows because there is no simple way to set such things as structs or arrays by typing values into an entry field). For the time being, setting the VarProbe to be non-interactive will not interfere with the drag&drop capability of objects into the *VarProbe* entry-field... This is considered an omission and will be corrected by the next release. Also, for the time being there is no way to disable a *VarProbe* once it has been displayed (i.e. after a *ProbeDisplay* containing the *VarProbe* has been created). The ability to do this will also be included in the next release.
- (const char *)getProbedVariable
  returns a string matching the identifier of the referent.
- (int)getDataOffset
  This is another example of a method used(?) exclusively by the Swarm kernel, it returns the offset used to extract the variable from the target object.

- (void *)probeRaw: anObject

- (void *)probeAsPointer: anObject

- (int)probeAsInt: anObject

- (double)probeAsDouble: anObject

- (const char *)probeAsString: anObject Buffer: (char *) buffer

- (const char *)probeAsString: anObject Buffer: (char *) buffer withPrecision: (int) precision

  These are the methods which actually use the probe. *probeRaw* simply returns a pointer to the probed variable, whereas the other methods return it as a particular type. *probeAsString* simply prints the value of the variable into the buffer (in more or less the same format you would expect from a *sprintf()*). The buffer should be pre-allocated. The second form of *probeAsString* is only used internally by ObjectSaver (../simtools/ObjectSaver.html) to use the "saved as" precision form which can be different from the "displayed" precision. See the section on setting the probing precision (../probing.html#precision) for floating-point numbers. Note: The *VarProbe* will complain if you try to do weird type conversions: for example, if the probed variable is of type pointer and you try to use the *probeAsDouble* method then the return value will be undefined (and you will get an error message on *stderr*).

- setStringReturnType: returnType

  When the *probedVariable* is of type unsigned char or char, the method *probeAsString* will, by default, return a string of the format: "'%c' %d". This is meant to reflect the commonplace use of an unsigned char as a small int. However, the user can tailor the format returned by the probe to use only one of the two possible representations.

The legal values for returnType are:

- DefaultString ("'%c' %i")
- CharString ("'%c'")
- IntString ("%i").

Note: When setting the value of an unsigned char or a char using the *setData:ToString:* (method described below), the expected format of the string is always "%i" unless *CharString* was chosen (in which case the format should be "'%c'").

- setData: anObject To: (void *)newValue

- setData: anObject ToString: (const char *)s

  The programmer can set the *probedVariable* via these methods, either by providing a pointer to the new value or a string which the probe will read and convert appropriately. Note: For the time being, the kernel itself is the only software which uses probes in this fashion, so we can get away with expecting either a string or a pointer. However, if any user should require more types (such as the choice provided in the *probeAs* series), please email (mailto:swarm@santafe.edu) us and I will include them in the following release!

(Probe.html)

(SwarmObject.html)

(MessageProbe.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

MessageProbe Class

```
//
// MessageProbe --
//   Allows the user to call a given message on any candidate
//   which is an instance of, or inherits from, a given class
//
@protocol MessageProbe <Probe>
CREATING
- setProbedSelector: (SEL) aSel;
- setProbedMessage: (const char *)aMessage;
- setHideResult: (int) val;
USING
-(int)isResultId;                 // I doubt that a user will
-(int)isArgumentId: (int)which;   // ever need these.

- (const char *)getProbedMessage;
- (int)getArgNum;
- (const char *)getArg: (int)which;
- (const char *)getArgName: (int)which;
- (int)getHideResult;

- setArg: (int)which To: (const char *)what;

- dynamicCallOn: target resultStorage: (const char **)result;
- dynamicCallOn: target;

- updateMethodCache: anObject;
- (int)intDynamicCallOn: target;
- (float)floatDynamicCallOn: target;
- (double)doubleDynamicCallOn: target;
@end
```

This is a specialized subclass of the abstract class Probe. It completes the specification of a probe whose referent is a message element of an object.

Create Phase Protocol

setProbedSelector: (SEL)aSelector

setProbedMessage: (const char *)aMessage
sets the message to be probed. At least one of these methods must be called.

Note that due to a bug in gcc-2.7.2 it is unadvisable (though we have not disallowed it) to make a message probe using setProbedMessage. The next release of gcc will contain a patch for this bug. If you absolutely require the ability to create messageprobes using setProbedMessage, please email (mailto: swarm@santafe.edu) us.

Usage Protocol

(const char *)getProbedMessage
: returns the string matching the identifier of the message being probed.

(int)getHideResult
: setHideResult: (int)val By setting hideResult to 1, the user is indicating that the result field, in a graphical representation of the message probe, should not be shown. In other words, if a probe display attempts to represent the message probe, it should hide the result window. This is useful, particularly when the return value is defined to be self.

(const char *)getArg: (int)which
: returns a string representation of the n'th argument.

(int)getArgNum
: returns the number of arguments the message takes.

(const char *)getArgName: (int)which
: returns a string representation of the argument with the given name.

setArg: (int)which To: (const char *)what
: sets the n'th argument of the message - the argument must be provided in string form.

(int)isResultId
: This method is used by the Swarm internals to find out whether the return value of the message is of type object.

(int)isArgumentId: (int)which
: This method is used by the Swarm internals to find out whether a given argument of the message is of type object.

updateMethodCache: anObject
: The message probe will always attempt to optimise the way in which it dispatches dynamic calls to the probed objects. In practice, this means that when the probed message takes no arguments, then the message probe will cache the actual implementation of the method and generate direct C function calls (as opposed to calling perform, for example) if the user promises that only one class of objects is being probed. The user does this by calling the updateMethodCache: with an example object of the appropriate target class. This method can be called as often as necessary throughout the lifetime of the message probe.

dynamicCallOn: target

dynamicCallOn: target resultStorage: (const char **)result

(int)intDynamicCallOn: target

(double)doubleDynamicCallOn: target

These are the methods used to actually generate a dynamic message call on the target object. The first method does not return a result, the second will return a dynamically malloc'ed string based representation of the result, whereas the third and fourth methods assume the user knows the exact typing of the method and would like a direct translation into the appropriate datatype.

(VarProbe.html)

(SwarmObject.html)

(ProbeMap.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

CustomProbeMap Class

```
//
// CustomProbeMap --
//   A subclass of ProbeMap, whose initial state is empty unlike the default
//   probeMap initial state which contains all the VarProbes
//   of the requested target class.
//
@protocol CustomProbeMap <ProbeMap, CREATABLE>
+ create: aZone forClass: (Class)aClass withIdenti-
fiers: (const char *)vars, ...;
@end
```

This subclass of the ProbeMap is used to create probe maps which are initialised in an emtpy state or with the VarProbes and MessageProbes intended. In other words, the probed class is set, as is the case with the normal ProbeMap class but upon createEnd no VarProbes or MessageProbes will be present within it. This feature is useful when creating a probe map from scratch (e.g. to be used in conjunction with the setProbeMap:For: message of the ProbeLibrary).

An additional convenience method was added to help compress the amount of code needed to create a custom probe.

**+ create: aZone forClass: (Class)aClass withIdenti-
fiers: (const char *)vars, ...;**

This is a variable argument message which takes a Zone, a Class name, and a colon-delimited list of instance variable names and message names valid for that class. The instance variable identifiers are listed first followed by a string consisting solely of a ":" followed by the message identifiers. The general format would be:

```
[CustomProbeMap create: aZone
              forClass: aClass
       withIdentifiers: "variable1", "variable2", ...,
                        ":",
                        "message1", "message2", ...,
                        NULL];
```

All strings up to the ":" are read as instance variable names and placed in the probe map and all strings after the ":" are read as message names and placed in the probe map. The NULL string terminates the list. An example of this message being used is given in the Advanced Usage Guide
(./AdvancedUsage.html)
.

(ProbeMap.html)

(Interface.html)

(CompleteProbeMap.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

CompleteProbeMap Class

```
//
// CompleteProbeMap --
//   A subclass of ProbeMap, whose initial state contains the VarProbes and
//   MessageProbes of the requested target class but also those of all its
//   superclasses.
//
@protocol CompleteProbeMap <ProbeMap> @end
```

Upon creation, this subclass of the ProbeMap will contain all the variables and all the messages of a given class (including the inherited ones).

(EmptyProbeMap.html)

(Interface.html)

(CompleteVarMap.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Objectbase Library: Usage Guide

Overview

The objectbase library contains the most basic objects users need to design their agents and swarms. It also serves, at present, as a repository for the probe machinery, which is provided for every SwarmObject. The way the classes in this library are to be used is varied. But, basically, it is provided so that the user will have something to subclass from for her own objects and Swarms.

Example Usage of SwarmObject

The best way to explain how the library should be used is to walk through an example. So, using Heatbugs, we'll walk through the ways objectbase is used and discuss them. Since more documentation is usually better than less, I'm going to explain things at a low level so that those not familiar with Objective C will understand the discussion. If you already are familiar with Objective C, then you should skip this part.

First off, the basic elements of the Heatbugs simulation are the heatbugs, the model swarm (which bundles the heatbugs), and the observer swarm (which bundles the displays of the probes poking into the model swarm and the heatbugs). The interface files for each show what must be imported and the declaration syntax needed to subclass from SwarmObject.

We'll use `Heatbug.h` for our discussion here. The first part of the file shows the C-preprocessor imports needed:

```
#import <objectbase/SwarmObject.h>
#import <space.h>
#import "HeatSpace.h"
#import <tkobjc/Raster.h>
```

The `#import <objectbase/SwarmObject.h>`; is included in order to subclass from SwarmObject. However, to provide backwards compatibility, we've placed this import in the library interface file `objectbase.h` as well, which means one could subclass from SwarmObject by simply importing the `objectbase.h` file. This is discouraged in order to make the library interfaces as standard as possible.

The next objectbase relevant piece of code in this file is:

```
@interface Heatbug: SwarmObject
{
  double unhappiness;
  int x, y;
  HeatValue idealTemperature;
  HeatValue outputHeat;
```

```
  float randomMoveProbability;

  Grid2d * world;
  int worldXSize, worldYSize;
  HeatSpace * heat;
  Color bugColor;
}
```

The @interface keyword indicates that you are beginning the definition of the part of an object (a Heatbug in this case) that will be visible to other objects. The Heatbug: SwarmObject indicates that you are calling this object Heatbug and it is a subclass of SwarmObject. What follows between the curly braces ({}) are the instance variables defined for the Heatbug class above and beyond those inherited from the SwarmObject class.

Inside this "agent," we have defined several parameters associated with either the agent, itself, or the space in which it sits. Any data that will need to be present throughout all the behavior and lifetime of the agent should be declared here. Also, anything declared here will be accessible to the probe machinery, and so will be capable of being manipulated and viewed from outside the agent.

Next come the message prototypes to which this agent will respond. And it is worth noting again that these are in addition to those declared in the SwarmObject superclass. So, not only will other objects be able to send messages to this agent that are declared here, but other objects will be able to send all the messages declared in the objectbase/SwarmObject.h imported previously. The messages prototyped here will dictate what the compiler thinks this object can respond to. Hence, if any part of any of these prototypes differs from the corresponding function definition in the Heatbug.m file, then the compiler will say something like Object: aHeatbug does not respond to xyz, where "xyz" is the name of the message that is being sent to the Heatbug. A script is provided with the Swarm distribution that fixes header file prototypes to match the message declarations in the corresponding ".m" file. This script should be in the $SWARMHOME/bin directory and is called m2h.

One more thing to notice about these prototypes is that some of them are duplicates of what appears in the objectbase/SwarmObject.h file. This means that when the message is called on a Heatbug object, it will execute the method defined here and not the one in the SwarmObject class. In the objectbase library, the following messages are intended to be overridden, as necessry: create:, createBegin:, createEnd, customizeBegin:, customizeEnd, customizeCopy:, describe:, and getInstanceName. Each of these messages do specific things that may change from subclass to subclass of SwarmObject. In this case, however, we're only overriding createEnd. The differences between we implement it in Heatbugs and the default is not that significant. But, it should be pointed out that when overriding certain messages, like createBegin: and createEnd, the new method should call the superclass' version of the message, as well. This is done using the default pointer to the superclass, designated super. The syntax in the Heatbugs case is:

```
[super createEnd];
```

The reasons for doing this are related to the object phase protocols used by defobj. If you would like more info on that, see Objective C - Swarm Style (../../swarmObjC.html) in the section entitled *The Create Phase In Principle*.

Finally, the `@end` keyword signifies the end of the interface definition. GNU Objective C allows one to leave this off; but, it is not good practice.

And that's it. Of course, there're a few tricky aspects to using the objectbase library that weren't mentioned here. Some of them will be mentioned in the Advanced Usage Notes (AdvancedUsage.html) and the Implementation Notes (Implementation.html) ; but, the best way to learn is to examine the way the demo applications do it and try to make some changes yourself.

Subclassing from Swarm

Subclassing from the Swarm class works very similar to subclassing from SwarmObject.

ActivityControl

The ActivityControl object provides much more finely grained control over the execution of an interactive simulation. It addresses both the problems of not being able to stop the simulation at any given point in any given activity and provides an initial step towards a Swarm debugger.

An activity controller can be attached to any activity that is created in a Swarm simulation, including those that are created for use only by the Swarm kernel. The controller then provides the basic activity manipulation messages on that activity, which are: `run`, `stop`, `next`, `step`, `stepUntil`, and `terminate`.

The presence of the ActivityControl object might cause some confusion about what role the ControlPanel should play in the controlled execution of the various schedules. The ControlPanel should still be used for the top-level control of any simulation that is running in a context where random interference is expected (like under a GUISwarm where the user may click a button at any time). The reason this is true is because the ControlPanel sends pseudo-interrupts to the infinite loop we use to perpetuate execution of the top level Swarm (which can only be seen in the form of the `go` message on a GUISwarm at present). This type of control may change in the future! But, for now, it is how we monitor the control state of the simulation.

Now, having said that, the ControlPanel should no longer be used to run the simulation.

It should only be used to instantiate the control context and quit the entire simulation. That means that sometime in the future, the Go and the Stop buttons will be removed from the ControlPanel display. They have been left in for backwards compatibility so that applications that do not use the new ActivityControl will retain their (albeit handicapped) controllability. Also, the current Time Step button will be renamed to Start to be consistent with it's new purpose.

In order to use the new control mechanism, you must place code like the following in the top-level Swarm. (This code was taken from a modified mousetrap demo app.)

```
observerActCont = [ActivityControl createBegin: [self getZone]];
  observerActCont = [observerActCont createEnd];
  [observerActCont attachToActivity: [self getSwarmActivity]];
  [probeDisplayManager createProbeDisplayFor: observerActCont];
```

This creates an ActivityControl and attaches it to the top-level activity (in this case an observerSwarm). It also creates a display for the controller. (The probe map for the ActivityControl class is designed within the ActivityControl, itself. This is done because all of these objects are expected to look the same to any outside object.) With this activity controller, you will then be able to run, stop, next, step, stepUntil, and terminate that activity.

There are some tricky aspects to successfully using an ActivityControl object that the Advanced Usage Notes
(AdvancedUsage.html)
will cover.

(Interface.html)

(../objectbase.html)

(AdvancedUsage.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Objectbase Library: Advanced Usage Notes

1. ProbeMap design: When designing a ProbeMap for a given (subclass of) SwarmObject, inclusion of instance variables or messages defined in the super class might be desirable. The normal ProbeMap design code might look like (this code was taken from the tutorial app called "hello-world"):

```
probeMap = [CustomProbeMap createBegin: [self getZone]];
      [probeMap setProbedClass: [person class]];
      probeMap = [probeMap createEnd];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "room"
 inClass: [person class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "party"
 inClass: [person class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "name"
 inClass: [person class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "stillhere"
 inClass: [person class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "listOfFriends"
 inClass: [person class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "myColor"
 inClass: [person class]]];
      [probeLibrary setProbeMap: probeMap For: [person class]];
      [probeDisplayManager createProbeDisplayFor: person];
```

where `room`, `party`, `name`, `stillhere`, `listOfFriends`, and `myColor` are instance variables declared in the interface to the Person subclass. And Person is a subclass of Agent2d, which is a subclass of SwarmObject.

Now let's say that there are two variables declared in Agent2d that you want to put into this custom probe in addition to the ones you've picked out of Person. Call them `x` and `y`. The way to add them to the `probeMap` is to add the following two lines of code to the above.

```
[probeMap addProbe: [probeLibrary getProbeForVariable: "x"
 inClass: [Agent2d class]]];
      [probeMap addProbe: [probeLibrary getProbeForVariable: "y"
 inClass: [Agent2d class]]];
```

And that's it! The two superclass-declared variables, which are, in fact, instance variables of the instance of the subclass, are now included in the probe.

In addition, a convenience message has been added to the CustomProbeMap interface to compress the above rather cluttered mechanism into one message. This convenience message can be used in the usual case where a ProbeMap will consist of variables and messages from the same class. For example, the first part of the custom probe creation above can be shortened to:

```
[probeLibrary create: [self getZone] forClass: [person class]
      withIdentifiers: "room", "party", "name", "stillhere",
```

```
                     "listOfFriends", "myColor", NULL];
```

And if the user wanted messages in the probe as well, it could be extended to:

```
probeMap = [CustomProbeMap create: [self getZone]
                             forClass: [person class]
                   withIdentifiers: "room", "party", "name",
                       "stillhere", "listOfFriends", "myColor",
                       ":",
                       "setWorld:Room:Party:",
                       "setPerson:Topic_array:ShowSpeech:",
                       NULL];
```

At present, this message doesn't search the superclasses for the message names listed here. But, that will soon be rectified.

2. ActivityControl Issues: It is completely reasonable to assume that explicit control can be had over all the activities in a simulation. However, at present, this control is limited because the context in which an activity runs determines how it behaves. To understand how an ActivityControl is to be used, we will have to explore the behavior of activities in context. (For a more complete explanation of the behavior of activities, see the activity library (../activity.html) documentation.)

   There are two ways to get an activity started, one can activate the activity in `nil` or in some other activity. So called "top-level" activities are activated in `nil` and are intended to be independent of the sophisticated scheduling activity that dictates the execution of actions in any other context in the simulation. I.e. the only activities that should be activated in `nil` are those sets of events that are expected to preserve the same behavior no matter what goes on in any other part of the simulation.

   The other type of activity, those activated in some other activity, is intended to be an integral part of its owner activity. However, this doesn't mean that it must depend on the outcome of actions happening in the owner activity. In fact, an ActionPlan can be designated as containing actions that are capable of being processed in parallel, via code like the following:

```
[anActionPlan setDefaultOrder: Concurrent];
```

But these activities are still intended to be meshed with their owner activities. In other words, they are part and parcel of the same model or simulation.

   Now, the operational effect of activating an activity in `nil` is that it will not be meshed with the rest of the Swarm activity structure. This gives the user (or process) complete control over the execution of that activity. A `run` on that activity will run either to completion or until a stop flag is set by a sequence of events purely internal to that activity. Or, one can stop it from the outside with a message from something like an ActivityControl.

   What all this means is that, while one can attach an ActivityControl to any activity, only the "top-level" activities (those having been activated in `nil`) are going to respond well to it. Any sub-activity will respond half-heartedly, if at all. For example, in the Mousetrap demo distributed with Swarm, an ActivityControl has been placed on both the ObserverSwarm and the ModelSwarm activities. Now,

if one sends a run message to the ActivityControl that is attached to the observerSwarm's activity, the entire model continues to run to completion, unless the user sends a stop message. However, if the sim is stopped at some point, a run message to the modelSwarm's activity will have no effect at all. (Note: If you do this via the activity controllers, you see the currentTime variable get updated; but, the actual run message to the activity, itself, has no effect.)
So, the rule of thumb, for the present, is to attach ActivityControl objects only to "top-level" activities, like the ObserverSwarm.

(UsageGuide.html)

(../objectbase.html)

(InterfaceDesign.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Swarmobject Library: Interface Design Notes

Unfortunately, this interface has not undergone a rigorous design review. As such it is subject to change in the future. However, there are rumours that this library will be integrated into the defobj library
(../defobj.html)
anyway. So, even though little thought was given to the design of this interface (and it is not likely to be worthwhile designing a robust interface at this time), it was implemented in order to provide a first step towards bringing all the various libraries in line with the standard set by defobj.

Along these lines, a few notes are relevant.

1. Probes may become an inherent part of any object.
2. ActivityControls will become a part of a larger set of tools used for debugging Swarm models.
3. The explicit importation of the SwarmObject interface should be removed. This is only provided for backward compatibility and when all the libraries follow the interface conventions, then those conventions should be encouraged.

(AdvancedUsage.html)

(../objectbase.html)

(Implementation.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Swarmobject Library: Implemenation Notes

1. Right now, Probe's rely on a special method, `getInstanceName`, that has to be implemented in any probe-able object in order to get anything other than the class name of that object into the object designation widget. However, a more general capability has been added to defobj to give a meaningful name to any object. Probes will be changed to take advantage of this new capability.
2. In the ActivityControl the frequency of the message `updateStateVar` is very high. It is sent at least once every cycle and every time a message is sent to the ActivityControl. This is unsatisfactory. Some of these messages can be pruned out of the object.
3. The `probeMap` designed for use with an ActivityControl was chosen fairly arbitrarily. Right now, it serves as a default for the class. A user can override it by designing a new one and inserting it into the `probeLibrary`.
4. Errors specific to objects in the objectbase library need to be gathered and initialized like those in the defobj library.

(InterfaceDesign.html)

(../objectbase.html)

(../objectbase.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Swarm Space Library

The Swarm Space library is the beginnings of a library to assist in building environments for interacting agents. In general, environments can be just as varied as the agents themselves (in one view, the environment itself is simply another agent). However, many simulations have similar types of environments that can be helpfully supported by generic code.

> The current space library only addresses simple kinds of discretized 2d space. Improvement is planned in the future: see the todo list
> (../todo.html)
> for ideas. Briefly, coordinates need to be elevated to the status of objects, which should hopefully allow spaces of different scales and boundary conditions to interact through a common reference system. In addition, other types of spaces are desired: continuous coordinates, other dimensions, arbitrary graphs, etc.

Discrete2d

Root class of all 2d discrete spaces. A Discrete2d is basically a 2d array of ids. Subclasses add particular space semantics onto this. Currently Discrete2d grids are accessed by integer pairs of X and Y coordinates.

Creation

- `setSizeX: (int)x Y: (int)y`
  Set the world size.
- `createEnd`
  Create the lattice, precompute the offsets based on Y coordinate.

Querying Discrete2d state

- `(int)getSizeX` and `- (int)getSizeY`
  Get the size of the lattice.

Global Operations

- `fastFillWithValue: (int)aValue`
- `fastFillWithObject: anObject`
- `fillWithValue: (int)aValue`

- fillWithObject: anObject

    These methods will fill the entire space with a specified value or object. The difference between the 'fill' and 'fastFill' methods is that the former use putObject and putValue internally, in other words, they are sensitive to any subclassing and alteration of these methods, whereas the latter use low level access routines
    directly and are therefore much faster (and remain unaffected by subclassing).

- copyDiscrete2d: (Discrete2d *)a toDiscrete2d: (Discrete2d *)b;

    This method copies the data in one Discrete2d object to another Discrete2d object. It assumes that both objects already exist.

- (int)setDiscrete2d: (Discrete2d *)a toFile: (const char *)filename;

    This method reads a PGM formatted file and pipes the data into a Discrete2d object.

Accessing elements

The Discrete2d is essentially a 2d array of ids. However, we also allow users to access the array as if it were an array of integers: the integers are cast into ids for storage. The Object methods retrieve values as objects, the Value methods cast things to integers before returning. We believe this is safe on all architectures.

- getObjectAtX: (int)x Y: (int)y

    Return the pointer stored at (x,y).

- getValueAtX: (int)x Y: (int)y

    Return the integer stored at (x,y).

- putObject: anObject atX: (int)x Y: (int)y

    Put the given pointer to (x,y) overwriting whatever was there.

- putValue: (int)v atX: (int)x Y: (int)y

    Put the given integer to (x,y) overwriting whatever was there.

Low level access

For speed, sometimes you want to go below the method call layer and access cells quickly. To allow this Discrete2d defines some macros that give you internal access. Use these at your own peril! In particular, subclasses have no way to redefine the macros.

- (id *)getLattice

    Returns the lattice pointer - use this for fast access.

- discrete2dSiteAt(lattice, offsets, x, y)

    Macro to return a pointer to the value at site x, y. lattice is the return value from getLattice, offsets is the precomputed offsets (stored in self->offsets).

Miscellaneous

These methods are used in `- createEnd` to actually allocate the array. Subclasses might need to use these.

- `- makeOffsets`
  Given an array size, compute the offsets array that caches the multiplication by ysize. See the `discrete2dSiteAt` macro.
- `- (id *)allocLattice`
  Allocate memory for the lattice.

Grid2d

Grid2d is a generic container class to represent agent position on a 2d lattice. It gets most of its behaviour from Discrete2d, adding extra code to check that you don't overwrite things by accident. Grid2d is pretty primitive: only one object can be stored at a site, no boundary conditions are implied, etc. A fancier Grid2d is in the works.

New Methods

- `- setOverwriteWarnings: (BOOL)b`
  If set to true, then if you try to store something at a site that doesn't have 0x0 there, a warning will be generated.

Overridden methods

- `- putObject: anObject atX: (int)x Y: (int)y`
  Replaces the Discrete2d method. First check to see if it should do overwrite warnings, and if so if you're going to overwrite: if both conditions are true, print out a warning message. Regardless of the check, it writes the new object in.
- `+ createBegin: aZone`
  Make overwrite warnings be on by default.

DblBuffer2d

DblBuffer2d augments Discrete2d to provide a form of double buffered space. Two lattices are maintained: lattice (the current state), and newLattice (the future state). All reads take place from lattice, all writes take place to newLattice. newLattice is copied to lattice when - updateLattice is called. DblBuffer2d can be used to implement one model of concurrent action, like in Ca2ds. NOTE: be very careful if you're using low-level macro access to the world, in particular be sure that you preserve the write semantics on the newLattice.

New Methods

- updateLattice
  Copy newLattice to lattice, in effect updating the lattice.
- (id *)getNewLattice
  Like - (id *)getLattice: return a pointer to the newLattice buffer.

Overridden Methods

- createEnd
  Rewrites the method from Discrete2d. Allocate two lattices, makes the offsets.
- putObject: anObject atX: (int)x Y: (int)y and - putValue: (int)value atX: (int)x Y: (int)y and
  Overridden so writes happen to newLattice.

Ca2d

Inherits from DblBuffer2d, defines abstract protocol for cellular automata.

New Methods

- setNumStates: (int)d
  Record the number of states the CA understands.
- initializeLattice
  Use this to set up your CA to a default initial state. Unimplemented in Ca2d.
- stepRule
  One iteration of the CA rule. Unimplemented in Ca2d.

Overridden Methods

- createEnd

Check that numStates has been set.

ConwayLife2d

Classic 2d Conway's Life CA.

Overridden Methods

+createBegin: aZone
: Set number of states to 2.

- initializeLattice
: Initialize lattice to random 1/3 in state 1.

- stepRule
: Run Conway's Life rule (simpleminded version).

Diffuse2d

Discrete 2nd order approximation to 2d diffusion with evaporation. Math is done in integers on the range [0,0x7fff].

New Methods

- setDiffusionConstant: (double)d
: Set the diffusion constant. Values over 1.0 might not be valid.

- setEvaporationRate: (double)d
: Set the evaporation rate. Values over 1.0 don't make much sense.

Overridden Methods

+createBegin: aZone
: Set diffusion constant and evaporation rate to 1.0, numStates to 0x7fff.

- initializeLattice
: Initialize world to 0.

- stepRule
: Run discrete approximation to diffusion. Roughly, it's

```
newHeat = evapRate * (self + diffuseConstant*(nbdavg - self))
where nbdavg is the weighted average of the 8 neighbours.
```

Object2dDisplay

Object2dDisplay helps display 2d arrays of objects. Create a Object2dDisplay, give it a Raster widget to draw on, a Discrete2d, a message to call on each object, and (optionally) a collection of objects and it will dispatch the message to all objects with the Raster widget as an argument. In addition, Object2dDisplay can help you make probees.

Creation Phase

- setDisplayWidget: (Raster *)r
  Set the display widget to use for drawing.
- setDiscrete2dToDisplay: (Discrete2d *)c
  Set the 2d array to draw
- setDisplayMessage: (SEL)s
  Set the message to be sent to each object in the grid to make it draw itself.
- setObjectCollection: objects
  (optional) set a collection of objects to be displayed. If this is not given, then Object2dDisplay loops through the 2d grid sending draw messages to all objects it finds there. Giving an explicit collection of objects to draw is more efficient if your grid is sparsely populated.

Use

- display
  Draw all objects in the array (or optionally, the collection) on the raster widget. All that happens here is the display message is sent to each object - it is the object's responsibility to render itself.
- makeProbeAtX: (int)x Y: (int)y
  Find an object at the given (x,y) coordinate and build a probe display for it. This is a good method to make a button client for a raster widget like

[aRaster setButtonClient: aObject2dDisplay Message: M(makeProbeAtX:Y:)]

Value2dDisplay

Value2dDisplay helps display 2d arrays of values. Value2dDisplay goes through a given Discrete2d array, turn states into colours, and draws them into a Raster widget.

Creation Phase

- setDisplayWidget: (Raster *)r Colormap: (XColormap *)c
  Set the display widget and the colourmap to use to draw the value array.
- setDiscrete2dToDisplay: (Discrete2d *)c
  Set which array to draw.
- setDisplayMappingM: (int)m C: (int)c
  Linear transform of states to colours for drawing.
  color = state / m + c
  If not set, assume m == 1 and c == 0.

Use

- display
  Draw the array on the given widget. Note that you still have to tell the widget to draw itself afterwards. The code for display uses the fast macro access in Discrete2d on the cached return value from getLattice. It also caches the drawPointX:Y: method lookup on the display widget - this is a nice trick that you might want to look at.

Int2dFiler

The Int2dFiler class is used to save the state of any Discrete2d object (or a subclass thereof) to a specified file.

A Note on Output Formats

The issue of what format(s) the class should know about is still open. For the purposes that I have encountered, simply storing the values in space-delimited line format:

```
value0_0 value0_1 ...
value1_0 value1_1 ...
...
...
```

has been sufficient. In particular, the output of this class has been used, in conjunction with the fragstats
(ftp://ftp.fsl.orst.edu/pub/fragstats.2.0/)
package, to apply landscape metrics to the output of a Swarm simulation. Nevertheless, I (Manor Askenazi
(mailto: manor@santafe.edu)
) would be very interested in hearing from end-users about any other (presumably more complex formats), as the need arises.

- setDiscrete2dToFile: (Discrete2d *)aSpace
  Set the target space to be filled. This message can be used more than once, but often it is useful to keep one Int2dFiler per space (e.g. when the space is saved multiple times).
- setValueMessage: (SEL)aSelector
  This message is optional. It is used when the target Discrete2d contains objects. By sending each object the message specified by the selector, the Int2dFiler is able to get from the object an integer representing its state, which it then writes to the file.
- setBackground: (int)aValue
  This message is optional. It is used when the target Discrete2d contains objects. If a particular location in the space has no resident object, the argument of this message is the value which gets writtent to the file. The default background value is 0.
- fileTo: (const char *)fileName
  When the Int2dFiler receives this message, it opens a file called *fileName*, stores the state of a pre-specified space into it, and then closes the file.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1997-12-17

Random Library

Documentation and Implementation Status

Note: The new random library does not work in the same way as the old one. This means that some applications that used the random library provided with the 1.0.2 release will be broken. However, porting these applications to the new random library will be fairly easy since large efforts were made to adhere to the standard set with the last version and some backwards compatibility hooks were incorporated.

Porting to Random 0.7

Overview

First, a few thoughts on pseudorandom number generation. It's hard to do right. There are many problems: the root cause, of course, is that computer algorithms themselves are not random. But there are also problems with defining "random", coming up with good tests for generators, and implementing algorithms correctly and efficiently. The history of pseudorandom number generation in simulation work is mostly embarassing. This library attempts to do a decent job of generating random numbers, as well as documenting how things work and what shortcomings there are. If you want to learn more about random number generation, the bibliography [See: $(SWARMHOME)/src/random/docs/SOURCES.for.0.7] has useful notes. Knuth is the main reference in this realm, but too old to describe most of the particular generators used here.

> The code here represents an effort to implement several efficient, reasonably safe generators. The algorithms come from reading the literature [See: $(SWARMHOME)/src/random/docs/SOURCES.for.0.7]: these algorithms have been implemented as accurately as possible and run through some simple tests. There is always a chance that some algorithm here is no good - there's also the chance that it is implemented incorrectly.

For best results, library users should test these generators themselves in some domain-specific way. One easy way to do this is to run an experiment twice: once with one class of generator (say, PMMLCG), and once with another (say, SWB). If the results differ radically, then you can suspect the generator. If they don't, well, the generator still might be faulty.

There is also a tarball of test programs for the random number library on the Swarm ftp site:
SwarmTests-0.7.tar.gz
(ftp://ftp2.santafe.edu/pub/swarm/RandomTests-0.7.tar.gz)

Documentation

The *random* library is defined and documented according to the Library Interface Conventions

(defobj/library.html)
. The following standard sections of documentation are available:

- Interface Reference
(random/Interface.html)
- Usage Guide
(random/UsageGuide.html)
- Advanced Usage Guide
(random/AdvancedUsage.html)

- Interface Design Notes
(random/InterfaceDesign.html)
- Implementation Notes
(random/Implementation.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Documentation and Implementation Status

This is version 0.7 of *Random*. It was donated by Sven Thommesen. Version 0.6 was a reimplementation of most of Nelson Minar's original random with many changes and a new interface. This version adds many more generators and distributions and changes the interface somewhat.

The implementation may be a little unstable. It hasn't been fully tested. Sven has also contributed a test suite which we will make available on the web site. But, as with any pseudo-random number generation library, results obtained from it should be examined very closely. But we are reasonably sure the generators and distributions have been correctly implemented.

Documentation

More documentation is available within the Swarm source distribution, itself [Ref: $(SWARMHOME)/src/random/docs/. These docs consist of:

WHATS.NEW.in.0.7
: Additions to Random since v. 0.6

WHATS.CHANGED.in.0.7
: Changes to Random since v. 0.6

WHATS.LEFT.TODO
: Improvements that may make it into v. 0.8

README.Generators.v07
: Usage Guide for generator objects

README.Distributions.v07
: Usage Guide for distribution objects

CHOOSING.A.GENERATOR
: Notes on choosing a generator for your sim

DEFAULT.GENERATORS
: Default generators invoked by distributions

SOURCES.for.0.7

Literature references for implemented objects

doc.generators.table
: Summary descriptive data for generators

doc.generators.maxvalue
: Values returned by [myGen getUnsignedSample]

doc.quality.generators
: Results of statistical and timing tests

doc.timings
: Timing results for generators and distributions

methods.SimpleGenerators.h
: Methods implemented by a typical Simple generator

methods.SplitGenerators.h
: Methods implemented by a typical Split generator

Anyone who attempts to generate random numbers by deterministic means is, of course, living in a state of sin. --Jon von Neuman

---

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Random Library: Porting to Random 0.7

The changes from version 0.6 to 0.7 consist mostly of a restructuring of the implementation of the library. However, there are some interface changes, as well, and the porting of applications written to work with the 0.6 version might require changes to accomodate the new interface. Also, it is advised that you take a look at the new interface, regardless of whether your application compiles and runs with the new random because the old interface is deprecated and will eventually be removed and because some of the functionality of the library changed such that the execution of your application may differ.

Minimalist Porting Notes:

What is described here is the minimum a normal random user would have to change to get her apps running with the 0.7. For complete documentation of the changes made to the random library, check the files:

$(SWARMHOME)/src/random/docs/WHATS.CHANGED.in.0.7
and
$(SWARMHOME)/src/random/docs/WHATS.NEW.in.0.7

1. The 'magic numbers' used by setState/getState to identify generators and distributions have changed. The sizes of the records returned by getState have also changed. Therefore version 0.7 generators will refuse to setStatea generator based on version 0.6 data.
2. The setState method has been renamed to setStateFrom, and getState to putStateInto.
3. The verifySelf methods have been exorcised. Their code has been moved to a separate test program.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

Random Library: Interface Reference

The *random* library uses the module definition conventions (../defobj/module.html) established by the *defobj* library (../defobj.html). It also depends on standard supertypes and classes defined by this library.

Dependencies

Following are the other header files imported by <random.h>:

```
#import <defobj.h>
#import <random/generators.h>
#import <random/distributions.h>
#import <random/RandomDefs.h>
#import <random/RandomVars.h>
```

The *defobj* library interface is included to provide the basic object support. RandomDefs.h contains some C preprocessor macros and typedefs used in the library.

Object Types (with links to reference pages)

Contents:

The following shows the protocol inheritance trees and object definitions for generators and distributions. The complete set of methods implemented by these objects have been subdivided into sub-protocols of methods that belong together logically. Just click on a (sub-)protocol name to see what methods it implements. (You may want to review the section on Protocols in the Objective-C book here!)

In the list below, any line that contains the word CREATABLE defines an object that you can use in your program. (This is part of the Swarm DefObj machinery.) In other words, while 'InternalState' is a normal protocol (a list of method definitions), the name 'ACGgen' refers to both a protocol and a class that implements that protocol. (You create an ACG generator by saying myGen = [ACGgen create: ....]). Similarly, 'GammaDist' defines both a protocol and a class that implements that protocol.

All generators and distributions ultimately inherit from SwarmObject.

The two files 'methods.SimpleGenerators.h' and 'methods.SplitGenerators.h' in directory /random/docs show the complete sets of methods implemented for the two most common types of generator.

Generators

- InternalState
(Prot.html#InternalState)
- SingleSeed
(Prot.html#SingleSeed)
- MultiSeed
(Prot.html#MultiSeed)
- Simple
(Prot.html#Simple)
- SimpleOut
(Prot.html#SimpleOut)
- SplitSingleSeed
(Prot.html#SplitSingleSeed)
- SplitMultiSeed
(Prot.html#SplitMultiSeed)
- Split
(Prot.html#Split)
- SplitOut
(Prot.html#SplitOut)
- SingleShortGenerator
(Prot.html#SingleShortGen)
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed>
- SingleLongGenerator
(Prot.html#SingleLongGen)
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>
- CombinedShortGenerator
(Prot.html#CombinedShortGen)
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>
- CombinedLongGenerator
(Prot.html#CombinedLongGen)
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>
- SingleShortSplitGenerator
(Prot.html#SingleShortSplitGen)
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed>
- SingleLongSplitGenerator

(Prot.html#SingleLongSplitGen)
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

- CombinedShortSplitGenerator
(Prot.html#CombinedShortSplitGen)
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

- CombinedLongSplitGenerator
(Prot.html#CombinedLongSplitGen)
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

- ACGgen
(ACGgen.html)
<SingleLongGenerator, CREATABLE>

- C2LCGXgen
(C2LCGXgen.html)
<CombinedShortSplitGenerator, CREATABLE>

- C2MRG3gen
(C2MRG3gen.html)
<CombinedLongGenerator, CREATABLE>

- C2TAUSgen
(C2TAUS1-3gen.html)
<CombinedShortGenerator, CREATABLE>

  - C2TAUS1gen
  (C2TAUS1-3gen.html#C2TAUS1gen)
  : C2TAUSgen <C2TAUSgen, CREATABLE>

  - C2TAUS2gen
  (C2TAUS1-3gen.html#C2TAUS2gen)
  : C2TAUSgen <C2TAUSgen, CREATABLE>

  - C2TAUS3gen
  (C2TAUS1-3gen.html#C2TAUS3gen)
  : C2TAUSgen <C2TAUSgen, CREATABLE>

- C3MWCgen
(C3MWCgen.html)
<CombinedLongGenerator, CREATABLE>

- C4LCGXgen
(C4LCGXgen.html)
<CombinedShortSplitGenerator, CREATABLE>

- LCGgen
(LCG1-3gen.html)
<SingleShortGenerator>

  - LCG1gen
  (LCG1-3gen.html#LCG1gen)
  : LCGgen <LCGgen, CREATABLE>

- LCG2gen
  (LCG1-3gen.html#LCG2gen)
  : LCGgen <LCGgen, CREATABLE>
- LCG3gen
  (LCG1-3gen.html#LCG3gen)
  : LCGgen <LCGgen, CREATABLE>

- MRGgen
(MRG5-7gen.html)
<SingleLongGenerator>

  - MRG5gen
    (MRG5-7gen.html#MRG5gen)
    : MRGgen <MRGgen, CREATABLE>
  - MRG6gen
    (MRG5-7gen.html#MRG6gen)
    : MRGgen <MRGgen, CREATABLE>
  - MRG7gen
    (MRG5-7gen.html#MRG7gen)
    : MRGgen <MRGgen, CREATABLE>

- MT19937gen
(MT19937gen.html)
<SingleLongGenerator, CREATABLE>
- MWCAgen
(MWCA-Bgen.html)
<SingleLongGenerator, CREATABLE>
- MWCBgen
(MWCA-Bgen.html#MWCBgen)
<SingleLongGenerator, CREATABLE>
- PMMLCGgen
(PMMLCG1-9gen.html)
<SingleShortGenerator>

  - PMMLCG1gen
    (PMMLCG1-9gen.html#PMMLCG1gen)
    : PMMLCGgen <PMMLCGgen, CREATABLE>
  - PMMLCG2gen
    (PMMLCG1-9gen.html#PMMLCG2gen)
    : PMMLCGgen <PMMLCGgen, CREATABLE>
  - PMMLCG3gen
    (PMMLCG1-9gen.html#PMMLCG3gen)
    : PMMLCGgen <PMMLCGgen, CREATABLE>
  - PMMLCG4gen

(PMMLCG1-9gen.html#PMMLCG4gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>

- PMMLCG5gen
(PMMLCG1-9gen.html#PMMLCG5gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>
- PMMLCG6gen
(PMMLCG1-9gen.html#PMMLCG6gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>
- PMMLCG7gen
(PMMLCG1-9gen.html#PMMLCG7gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>
- PMMLCG8gen
(PMMLCG1-9gen.html#PMMLCG8gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>
- PMMLCG9gen
(PMMLCG1-9gen.html#PMMLCG9gen)
: PMMLCGgen <PMMLCGgen, CREATABLE>

- PSWBgen
(PSWBgen.html)
<SingleLongGenerator, CREATABLE>
- RWC2gen
(RWC28gen.html)
<SingleLongGenerator, CREATABLE>
- RWC8gen
(RWC28gen.html#RWC8gen)
<CombinedLongGenerator, CREATABLE>
- SCGgen
(SCGgen.html)
<SingleLongGenerator, CREATABLE>
- SWBgen
(SWB1-3gen.html)
<SingleLongGenerator>
  - SWB1gen
  (SWB1-3gen.html#SWB1gen)
  : SWBgen <SWBgen, CREATABLE>
  - SWB2gen
  (SWB1-3gen.html#SWB2gen)
  : SWBgen <SWBgen, CREATABLE>
  - SWB3gen
  (SWB1-3gen.html#SWB3gen)
  : SWBgen <SWBgen, CREATABLE>

- TGFSRgen
(TGFSRgen.html)
<SingleLongGenerator>

  - TT403gen
  (TGFSRgen.html#TT403gen)
  : TGFSRgen <TGFSRgen, CREATABLE>
  - TT775gen
  (TGFSRgen.html#TT775gen)
  : TGFSRgen <TGFSRgen, CREATABLE>
  - TT800gen
  (TGFSRgen.html#TT800gen)
  : TGFSRgen <TGFSRgen, CREATABLE>

Distributions

- ProbabilityDistribution
(Prot.html#ProbabilityDist)
<Create, Drop, InternalState>
- BooleanDistribution
(Prot.html#BooleanDistribution)
<ProbabilityDistribution>
- IntegerDistribution
(Prot.html#IntegerDistribution)
<ProbabilityDistribution>
- UnsignedDistribution
(Prot.html#UnsignedDistribution)
<ProbabilityDistribution>
- DoubleDistribution
(Prot.html#DoubleDistribution)
<ProbabilityDistribution>
- Normal
(Prot.html#Normal)
<DoubleDistribution>
-
  RandomBitDist
  (RandomBitDist.html)
  <BooleanDistribution, CREATABLE>
- BernoulliDist
(BernoulliDist.html)
<BooleanDistribution, CREATABLE>

- UniformIntegerDist
(UniformIntegerDist.html)
<IntegerDistribution, CREATABLE>
- UniformUnsignedDist
(UniformUnsignedDist.html)
<UnsignedDistribution, CREATABLE>
- UniformDoubleDist
(UniformDoubleDist.html)
<DoubleDistribution, CREATABLE>
- ExponentialDist
(ExponentialDist.html)
<DoubleDistribution, CREATABLE>
- GammaDist
(GammaDist.html)
<DoubleDistribution, CREATABLE>
- NormalDist
(NormalDist.html)
<Normal, CREATABLE>
- LogNormalDist
(NormalDist.html#LogNormalDist)
<Normal, CREATABLE>

(../random.html)

(../random.html)

(UsageGuide.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Protocol Reference

InternalState

```
//
// InternalState --
//   Methods to save the internal state of an object (genera-
tor, distribution)
//   to a memory buffer allocated by the calling pro-
gram, and to set the state
//   of an object from previously saved state data, provided in a mem-
ory buffer.
//   Method -getStateSize specifies the minimum buffer size needed.
//
@protocol InternalState
- (unsigned) getStateSize; // size of buffer needed
- (void) putStateInto: (void *) buffer; // save state data for later use
- (void) setStateFrom: (void *) buffer; // set state from saved data
- (void) describe: (id) outStream; // prints ascii data to stream
- (char *)      getName; // returns name of object
- (unsigned) getMagic; // object's 'magic number'
@end
```

The InternalState protocol provides attributes common to the generators and distributions provided in the random library. This protocol could be used for any class of objects that might contain an internal state, including things unrelated to pseudo random number generation.

The "state" of a given object depends on the type of object it is. In the case of random, both the bit generators and the distributions are capable of carrying state. The state is, in all cases, a structure that is understood by the object with which it is associated. The details of this structure can be found in the code for the given generator or distribution; but, in general, the user is not expected to understand the structure. Where the state must be saved and used to restart, the user is expected only to read and write the block of memory from and to the disk and pass the pointer to and from the object with which the state is associated. [See the Usage Guide (./UsageGuide.html) for examples of how to do this.] Please note that different computer architectures store binary data in memory in different ways (using different byte orders). Since this version of the InternalState protocol stores data in binary form, you cannot expect to be able to use these methods to transfer data between different computers. We hope to lift this restriction in a future version.

(unsigned) -getStateSize returns an unsigned integer indicating how many bytes of memory are necessary to save the state of the object.

(void) -getState: (void *) state writes a copy of the state of the object into the memory address pointed to by the "state" pointer passed to it. The memory should be allocated before the message is sent.

(void) -setState: (void *) state copies the values contained in the memory address pointed to by "state" into the state variables used by the object. The block of memory containing the state may be sent back to the heap after the message has been sent.

(./Interface.html)

(./Interface.html)

These descriptions are not finished. Please see the documentation in $(SWARMHOME)/src/random/docs/

Generator Protocols

SingleSeed

```
@deftype SingleSeed
CREATING
+  create: aZone setStateFromSeed: (unsigned) seed;
USING
- setStateFromSeed:  (unsigned)   seed;
- (unsigned) getMaxSeedValue;
- (unsigned)  getInitialSeed;
@end
```

MultiSeed

```
@deftype MultiSeed
CREATING
+ create: aZone setStateFromSeeds: (unsigned *) seeds;
```

```
USING
- setStateFromSeeds: (unsigned *) seeds;
- (unsigned)  lengthOfSeedVector;
- (unsigned *) getMaxSeedValues;
- (unsigned *)  getInitialSeeds;
@end
```

Simple

```
@deftype Simple
CREATING
+ createWithDefaults: aZone;
USING
- setAntithetic: (BOOL) antiT;
- (BOOL) getAntithetic;
- (unsigned long long int) getCurrentCount;
@end
```

SimpleOut

```
@deftype SimpleOut
USING
- (unsigned)    getUnsignedMax;
- (unsigned) getUnsignedSample;
- (float)       getFloatSample; // using 1 unsigned
- (double)      getThinDoubleSample; // using 1 unsigned
- (double)      getDoubleSample; // using 2 unsigneds
- (long double) getLongDoubleSample; // using 2 unsigneds
@end
```

SplitSingleSeed

```
@deftype SplitSingleSeed
CREATING
+ create: aZone
   setA: (unsigned) A  // # of virtual generators
   setv: (unsigned) v  // log2(#segments/generator)
   setw: (unsigned) w // log2(segment length)
setStateFromSeed: (unsigned) seed;
USING
- setStateFromSeed: (unsigned) seed;
- (unsigned) getMaxSeedValue;
- (unsigned) getInitialSeed;
@end
```

SplitMultiSeed

```
@deftype SplitMultiSeed
CREATING
+ create: aZone
   setA: (unsigned) A  // # of virtual generators
   setv: (unsigned) v  // log2(#segments/generator)
```

```
   setw: (unsigned) w // log2(segment length)
setStateFromSeeds: (unsigned *) seeds;
USING
- setStateFromSeeds: (unsigned *) seeds;
- (unsigned)  lengthOfSeedVector;
- (unsigned *) getMaxSeedValues; // min is 1
- (unsigned *) getInitialSeeds; // = getInitialSeeds: 0
@end
```

Split

```
@deftype Split
CREATING
+ createWithDefaults: aZone;
USING
- setAntithetic: (BOOL) antiT;
- (unsigned)    getNumGenerators; // returns A
- (unsigned)    getNumSegments; // returns v
- (unsigned)    getSegmentLength; // returns w
- (BOOL) getAntithetic;
- initGenerator:    (unsigned) vGen; // to segment #0
- restartGenerator: (unsigned) vGen; // start of current segment
- advanceGenerator: (unsigned) vGen; // to next segment
- jumpGenerator:    (unsigned) vGen  toSegment: (unsigned long long int) seg;
- initAll; // to segment #0
- restartAll; // start of current segment
- advanceAll; // to next segment
- jumpAllToSegment: (unsigned long long int) seg;
- (unsigned long long int) getCurrentCount:   (unsigned) vGen;
- (unsigned long long int) getCurrentSegment: (unsigned) vGen;
@end
```

SplitOut

```
@deftype SplitOut
USING
- (unsigned)    getUnsignedMax;
- (unsigned) getUnsignedSample:   (unsigned) vGen;
- (float)       getFloatSample:      (unsigned) vGen; // using 1 unsigned
- (double)      getThinDoubleSample: (unsigned) vGen; // using 1 unsigned
- (double)      getDoubleSample:     (unsigned) vGen; // using 2 unsigneds
- (long double) getLongDoubleSample: (unsigned) vGen; // using 2 unsigneds
@end
```

SingleShortGenerator
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed>

No extra messages added.

SingleLongGenerator
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>

No extra messages added.

CombinedShortGenerator
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>

No extra messages added.

CombinedLongGenerator
<Create, Drop, InternalState, SimpleOut, Simple, SingleSeed, MultiSeed>

No extra messages added.

SingleShortSplitGenerator
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed>

No extra messages added.

SingleLongSplitGenerator
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

No extra messages added.

CombinedShortSplitGenerator
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

No extra messages added.

CombinedLongSplitGenerator
<Create, Drop, InternalState, SplitOut, Split, SplitSingleSeed, SplitMultiSeed>

No extra messages added.

Distribution Protocols

ProbabilityDistribution
<Create, Drop, InternalState>

```
// ProbabilityDistribution --
//
// A process for generating a sequence of random numbers matching the
//   frequencies defined by a specific distribution function.  The
//   process is driven by input from a supplied uniform random generator.
//
@deftype ProbabilityDistribution <Create, Drop, InternalState>
CREATING
+ createWithDefaults: aZone;
+ create: aZone setGenerator: (id) generator;
+ create: aZone setGenerator: (id) generator
setVirtualGenerator: (unsigned) vGen;
+ createBegin;
USING
```

```
- setGenerator: (id) generator;
- setGenerator: (id) generator
setVirtualGenerator: (unsigned) vGen;
- createEnd;
- reset;
-               getGenerator;
- (unsigned) getVirtualGenerator;
- (BOOL) getOptionsInitialized;
- (unsigned long long int) getCurrentCount;
@end
```

BooleanDistribution
<ProbabilityDistribution>

```
//
// BooleanDistribution --
//   A probability distribution that returns YES/NO sample values.
//
@deftype BooleanDistribution <ProbabilityDistribution>
- (BOOL) getBooleanSample;
- (int) getIntegerSample; // for convenience
@end
```

IntegerDistribution
<ProbabilityDistribution>

```
//
// IntegerDistribution --
//   A probability distribution that returns integer sample values.
//
@deftype IntegerDistribution <ProbabilityDistribution>
- (int) getIntegerSample;
@end
```

UnsignedDistribution
<ProbabilityDistribution>

```
//
// UnsignedDistribution --
//   A probability distribution that returns non-
negative integer sample values.
//
@deftype UnsignedDistribution <ProbabilityDistribution>
- (unsigned) getUnsignedSample;
@end
```

DoubleDistribution
<ProbabilityDistribution>

```
//
// DoubleDistribution --
//   A probability distribution that returns an approximation of continuous
```

```
//   values as represented by double-precision floating point values.
//
@deftype DoubleDistribution <ProbabilityDistribution>
- (double) getDoubleSample;
@end
```

Normal
<DoubleDistribution>

```
//
// NormalDistribution -- a well-known continuous probability distribution
// LogNormalDistribution -- also well-known ...
//
@deftype Normal <DoubleDistribution>
CREATING
+ create: aZone setGenerator: generator
                setMean: (double) mean setVariance: (double) variance;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setMean: (double) mean setVariance: (double) variance;
USING
- setMean: (double) mean setVariance: (double) variance;
- (double) getMean;
- (double) getVariance;
- (double) getStdDev;
- (double) getSampleWithMean: (double) mean
withVariance: (double) variance;
@end
```

(./Interface.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

ACGgen Class

```
//
// ACGgen--
//   predefined random number generators, not recommended for serious use;
//   included for historical reasons (compatibility with earlier releases).
//
@protocol ACGgen      <SingleLongGenerator, CREATABLE> @end
```

Use of this generator is NOT Recommended

Documentation

```
Generator name:  | ACG

Description:  | Additive Congruential Generator

Reference:  | ACG is an often used algorithm, mentioned in Knuth,
| Numerical Recipes, etc.

| state vector = f(seed)
Algorithm:  | state_n = (state_(n-s) + state_(n-r)) mod 2^32; r > s
| output = state_n

| ACG is in the Lagged Fibonacci class of generators.
| These generators use a basic algorithm of the form
| X_n = f(X_(n-r),X_(n-s)) mod m; r>s
| The function f is typically xor, addition, subtraction,
| multiplication or subtraction with carry. It uses
| simpler math than a basic LCG, but keeps a larger state.

Comments:  | This generator relies on controlled overflow at 32 bits.
| This requires that unsigned be a 32bit value that follows
| ANSI C rules.

Implementation:  | The state vector contains the generator's r most recent
| outputs (X_n .. X_(n-r)). This vector is initialized from
| an LCG generator using one 32-bit seed value. After
| initialization we throw away r*5 numbers to 'prime' the
| generator (i.e. make sure we have entered a closed cycle.)
```

```
Parameters:  | unsigned int r // long lag
| unsigned int s // short lag

State:  | unsigned int state[r] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,2^32-1] may be returned
| (i.e. [0,4294967295].)

Valid seeds:  | Using create:setStateFromSeed:  [1,2^32-1]
| Using create:setStateFromSeeds: [1,2^32-2]

Cycles:  | With properly chosen parameters, these generators
| have a cycle of length about 2^r.

Output Quality:  | Low order bits are poor, and there are other problems.
| SWB is a better generator in the same class, and it is
| only slightly slower.

Reference for | Knuth, op.cit.
output quality:  |

ACGgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                      2.702 uS
|   -getFloatSample, getDoubleSample:         4.482 uS
|   -getFatDoubleSample, getLongDoubleSample: 8.299 uS

Relative speed:
| Speed 1.369 (time 0.731) relative to MT19937 getUnsignedSample
--------------- | -----------------------------------------------------------
---

Parameters:  | r = 55
| s = 24

Reference for | These values recommended by Knuth.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length ~ 2^r.

Output Quality:  | Low-order bits are poor.
```

Reference for | Knuth, op.cit.
output quality: |

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

C2LCGXgen Class

```
//
// C2LCGX --
//   short component based generator with splitting facilities. Recommended.
//
@deftype C2LCGXgen  <CombinedShortSplitGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | C2LCGX

Description:  | Combined random generator using 2 (PMM)LGC generators

Reference:  | Pierre L'Ecuyer and Serge Cote,
| "Implementing a Random Number Package with Splitting
| Facilities." ACM Transactions on Mathematical Software
| vol. 17 no. 1, March 1991, pp. 98-111.
| Their name for the generator: 'split.c'.
| See also <http://www.iro.umontreal.ca/~lecuyer>.

Comments:  | This portable generator is based on a backbone generator
| which is a combination of 2 (PMM)LCG generators. It has a
| period length of almost 2^61 (2.3e18).

| The backbone generator's period can be split up into a
| number of 'virtual generators' (A), each of which can be set
| to access a number of 'segments' (V) of length W, subject to
| the constraint that A * V * W <= 2^60.

| This is possible because the backbone generator has
| 'jump-ahead facilities' -- i.e. a quick method of calculating
| the state of the generator at a given point in its cycle
| without having to actually draw a large number of variates.

Algorithm | state(j) = seed(j)
for each | newstate(j) = (a(j) * oldstate(j)) mod m(j)
generator:  | output(j) = state(j)

Algorithm | output = (state(1) - state(2)) mod m(0)
for combination:|
```

```
Comments:  | For a discussion of PMMLCG generators, see the source file
| PMMLCGgen.h.

Configuration | unsigned int numGenerators; // # virtual generators
parameters:  | unsigned int numSegments; // log2(#segments/vGen)
| unsigned int segmentLength; // log2(segment length)

State:  | unsigned int state [numGenerators] [2];
| (each virtual generator has a 4-integer state vector)

Output types | unsigned int in [0,unsignedMax].
(for each | float in [0.0,1.0) using 1 iteration of the generator.
virtual | double in [0.0,1.0) using 1 or 2 iterations of the generator.
generator):  | long double in [0.0,1.0) using 2 iterations of the generator.

Output range:  | The combined generator yields output in the range [0,m0-1]
| (i.e. [0,2147483562].)

Limits on | Number of generators A <= 2^25  (33,554,432)
parameters:  | (You'll likely run out of memory before then ...)
| Number of segments / generator <= 2^60 (i.e. v <= 60)
| Length of each segment <= 2^60 (i.e. w <= 60)

Defaults:  | A=32, v=20, w=30

Valid seeds:  | Using create:setStateFromSeed: in [1, 2^32-1]
| Using create:setStateFromSeeds:
| The generator must be seeded with 2 starting seeds:
|  Seed[0]: [1, 2147483562] = [1,m1-1]
|  Seed[1]: [1, 2147483398] = [1,m2-1]

Cycles:  | With properly chosen parameters, the component generators
| have a single full cycle of length (m-2).
| The combined generator has a cycle length of
| (m1-1)*(m2-1)/2 = 2147483562*2147483398/2 ~= 2^61.

Output Quality:  | All bits should be safe.
| The combined generator has much better statistical
| properties than either of its component PMMLCG's.

Reference for | L'Ecuyer and Cote, op. cit.
output quality:  |

Implementation  | The code for -getUnsignedMax uses a trick devised by
comment:        | Schrage to avoid 32-bit overflow when calculating the
| next state (i.e. using q=m/a and r=m%a as auxiliary
| values). Using 64-bit math (long long integers) is more
```

| straightforward. Unfortunately, on x86 at least, there is
| a significant speed penalty for doing so.
| The output, however, would be identical. YMMV.

C2LCGXgen | On a 486/66, the following approximate measures were obtained:
speed: | -getUnsignedSample: 7.029 uS
| -getFloatSample, getDoubleSample: 9.049 uS
| -getFatDoubleSample, getLongDoubleSample: 16.725 uS

Relative speed:
| Speed 0.526 (time 1.900) relative to MT19937 getUnsignedSample
--------------- . ---
--

Parameters: | a[COMPONENTS] = { 40014, 40692 }
| m[COMPONENTS] = { 2147483563, 2147483399 }

Reference for | L'Ecuyer and Cote, op. cit.
parameters: |

Cycles: | With the chosen parameters, this generator
| has a single full cycle of length close to 2^60 (2.3e18).

Output Quality: | The authors performed an exhaustive search for component
| generators that would yield the best statistical quality.

Reference for | L'Ecuyer and Cote, op. cit.
output quality: |

(./ACGgen.html)

(./Interface.html)

(./C2MRG3gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

C2MRG3gen Class

```
//
// C2MRG3 --
//   long component based generator recommended for use.
//
@deftype C2MRG3gen <CombinedLongGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | C2MRG3

Description:  | Combination of 2 Multiple Recursive LCG Generators

Reference:  | Pierre L'Ecuyer,
| "Combined Multiple Recursive Random Number Generators."
| See <http://www.iro.umontreal.ca/~lecuyer>.

| For a discussion of MRG generators, see MRGgen.h.

Algorithm:  | For each generator:
| x_n = (a_1*x_n-1 + a_2*x_n-2 ... + a_k*x_n-k) mod m
| Combining them:
| x = (x1 - x2) mod m1

Comments:  | Combinations of like generators are shown to have better
| statistical properties than single generators.

Implementation:  | The components of this generator each has two nonzero
| multipliers (and one that's zero). They use different
| moduli (2^31-1, 2145483479.)

Parameters:  | unsigned int m[2] // modulus
| unsigned int a1[2], ak[2] // multipliers
| unsigned int k[2] // the order of the generator

State:  | unsigned int x[k][2] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
```

```
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,m1-2].

Valid seeds:  | This generator requires 3 seeds for each component:
| within [1,m1-1] for component 1;
| within [1,m2-1] for component 2.
| Using create:setStateFromSeed: in [1,2^32-1]

Cycles:  | With properly chosen parameters, this generator
| has a single cycle of length (m1^3-1)(m2^3-1)/2
| or 2^184 < cycle < 2^185 (4.9e55).

Output Quality:  | Good.

Reference for | L'Ecuyer et al, op. cit.
output quality:  |

C2MRG3gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        13.459 uS
|   -getFloatSample, getDoubleSample:         15.203 uS
|   -getFatDoubleSample, getLongDoubleSample: 29.699 uS

Relative speed:
| Speed 0.275 (time 3.640) relative to MT19937 getUnsignedSample
-------------- . ----------------------------------------------------------
--

Parameters:  | unsigned int k = 3 // lag order

| unsigned int m1 = 2^31 - 1 // modulus
| unsigned int a12 = 63308 // multiplier
| unsigned int a13 = -183326 // multiplier

| unsigned int m2 = 2145483479 // modulus
| unsigned int a21 = 86098 // multiplier
| unsigned int a23 = -539608 // multiplier

Reference for | L'Ecuyer et al, op. cit., Table 3.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length
| ((2^31-1)^3 - 1) * (2145483479^3 - 1) / 2 ~= 4.89e55,
| i.e. 2^184 < cycle < 2^185.
| (The paper, in apparent error, claims
| that the cycle is close to 2^205.)
```

```
Output Quality:  | Good. Parameters chosen on the basis of theory
                 | and empirical testing.

Reference for    | L'Ecuyer et al, op. cit.
output quality:  |
```

(./C2LCGXgen.html)

(./Interface.html)

(./C2TAUS1-3gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

C2TAUSgen Class

```
//
// C2TAUS1, C2TAUS2, C2TAUS3 --
//   short component based generator recommended for use.
//
@deftype C2TAUSgen    <CombinedShortGenerator> @end
@deftype C2TAUS1gen   <C2TAUSgen, CREATABLE> @end
@deftype C2TAUS2gen   <C2TAUSgen, CREATABLE> @end
@deftype C2TAUS3gen   <C2TAUSgen, CREATABLE> @end
```

General Documentation

```
Generator name:  | C2TAUS

Description:  | Combined Tausworthe Generator
| based on 2 component generators of periods 2^31-1 and 2^29-1

Reference:  | Shu Tezuka and Pierre L'Ecuyer,
| "Efficient and Portable Combined Tausworthe
| Random Number Generators."
| ACM Transactions on Modeling and Computer Simulation,
| vol. 1 no. 2, April 1991, pp. 99-112.

| Original reference for the Tausworthe algorithm:
| R. C. Tausworthe, "Random Numbers Generated by Linear
| Recurrence modulo 2", Math. Comput. 19 (1965), 201-209.

Algorithm:  | state = seed, then:

(Each | b = ((old << Q) ^ old) & mask
component:)  | new = ((old << S) ^ (b >> (P-S)) & mask

(Combination:)  | output = ((new1 ^ (new2 << (P1-P2))

Implementation:  |

Parameters:  | int   P [2] // word length
| int   Q [2] //
| int   S [2] //
| int PmS [2] // P - S
| int P1mP2 // P[0] - P[1]
```

```
| unsigned int Mask[2] // 2^P - 1

State:  | unsigned int state[2] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,2^31-1] may be returned
| (i.e. [0,2147483647].)

Valid seeds:  | Using create:setStateFromSeed: in [1, 2^32-1]
| Using create:setStateFromSeeds:
|  Seed[0] must be in [1,Mask[0]];
|  Seed[1] must be in [1,Mask[1]];

Cycles:  | These generators have a single full cycle
| of length Mask1 * Mask2 ~= 2^60 (1.15e18).

Output Quality:  | The authors conducted an exhaustive search among candidate
| Tausworthe generators to find the combination that had the
| best lattice structure. (Single Tausworthe generators do not
| have acceptable properties.) The 3 generators implemented
| here were retained as the best found.

Reference for | Tezuka and L'Ecuyer, op. cit.
output quality:  |

C2TAUSgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        4.078 uS
|   -getFloatSample, getDoubleSample:          5.932 uS
|   -getFatDoubleSample, getLongDoubleSample: 10.920 uS

Relative speed:
| Speed 0.907 (time 1.103) relative to MT19937 getUnsignedSample
```

C2TAUS1gen Documentation

```
Parameters:  | P = 31 so Mask = 2^31-1
(component 1) | S = 12 so PmS = 19
| Q = 13

| P = 29 so Mask = 2^29-1
(component 2) | S = 17 so PmS = 12
```

```
| Q =  2

Reference for | These values recommended by Tezuka and L'Ecuyer,
parameters:  | op. cit., table I p. 107.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length ~ 2^60.

Output Quality:  | Good.

Reference for | Tezuka and L'Ecuyer, op.cit.
output quality:  |
```

C2TAUS2gen Documentation

```
Parameters:  | P = 31 so Mask = 2^31-1
(component 1) | S = 21 so PmS = 10
| Q =  3

| P = 29 so Mask = 2^29-1
(component 2) | S = 17 so PmS = 12
| Q =  2

Reference for | These values recommended by Tezuka and L'Ecuyer,
parameters:  | op. cit., table II p. 107.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length ~ 2^60.

Output Quality:  | Good.

Reference for | Tezuka and L'Ecuyer, op.cit.
output quality:  |
```

C2TAUS3gen Documentation

```
Parameters:  | P = 31 so Mask = 2^31-1
(component 1) | S = 13 so PmS = 18
| Q = 13

| P = 29 so Mask = 2^29-1
(component 2) | S = 20 so PmS =  9
| Q =  2

Reference for | These values recommended by Tezuka and L'Ecuyer,
```

parameters: | op. cit., table III p. 107.

Cycles: | With the chosen parameters, this generator
| has a single full cycle of length ~ 2^60.

Output Quality: | Good.

Reference for | Tezuka and L'Ecuyer, op.cit.
output quality: |

(./C2MRG3gen.html)

(./Interface.html)

(./C3MWCgen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

C3MWCgen Class

```
//
// C3MWC --
//   long component based generator recommended for use.
//
@deftype C3MWCgen   <CombinedLongGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | C3MWC

Description:  | Combined Multiply-With-Carry generator

Reference:  | Article posted to Usenet in 1994 by George Marsaglia,
| "The Mother of all Random Number Generators."
| Source found at <ftp://ftp.taygeta.com/pub/c/mother.c>
| Also document /papers/mwc1.ps on his Diehard CD-ROM,
| found at <http://www.hku.hk/internet/randomCD.html>.

Algorithm:  | K = a * (K & 65535) + (K>>16);
| J = b * (J & 65535) + (J>>16);
| I = c * (I & 65535) + (I>>16);
| L = d * (L & 65535) + (L>>16);
| M = e * (M & 65535) + (M>>16);
| N = f * (N & 65535) + (N>>16);
| return ((K+I+M)>>16) + (J+L+N);

Comments:  | This generator is a combination of 3 MWC generators,
| each of which is a combination of 2 16-bit
| Multiply-With-Carry generators.

Parameters:  | unsigned int a = 30903 // multiplier
| unsigned int b = 18000
| unsigned int c = 29013
| unsigned int d = 30345
| unsigned int e = 30903
| unsigned int f = 31083

State:  | unsigned int K,J,I,L,M,N // state vector

Output types:  | unsigned int (we return the current state).
```

```
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,2^32-1].

Valid seeds:  | Using create:setStateFromSeed:  [1, 2^32-1]
| Using create:setStateFromSeeds: each component may have
| a 'state' <= 0xffff and a carry <= (multiplier-1), so that
| for the first component the largest seed allowed is
| (a-1)<<16 + 65535.

Cycles:  | This generator is claimed to be strictly periodic, with a
| period > 2^118. There may be more than one cycle.

Output Quality:  | Good.

Reference for | Marsaglia, op. cit.
output quality:  |

C3MWCgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        6.437 uS
|   -getFloatSample, getDoubleSample:          8.200 uS
|   -getFatDoubleSample, getLongDoubleSample: 15.562 uS

Relative speed:
| Speed 0.574 (time 1.741) relative to MT19937 getUnsignedSample
--------------- . -----------------------------------------------------------
--
```

(./C2TAUS1-3gen.html)

(./Interface.html)

(./C4LCGXgen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

C4LCGXgen Class

```
//
// C4LCGXgen --
//   short component based generator with splitting facilities. Recommended.
//
@deftype C4LCGXgen  <CombinedShortSplitGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | C4LCGX

Description:  | Combined random generator using 4 (PMM)LGC generators

Reference:  | Pierre L'Ecuyer and Terry H. Andres,
| "A Random Number Generator Based on the Combination
|  of Four LCGs", 17 December 1996.
| To appear in Mathematics and Computers in Simulation. See:
| <http://www.iro.umontreal.ca/~lecuyer/myftp/papers/clcg4.ps>.
| Their name for the generator: 'clcg4'.

Comments:  | This portable generator is based on a backbone generator
| which is a combination of 4 (PMM)LCG generators. It has a
| period length of (m1-1)(m2-1)(m3-1)(m4-1) / 8, or almost
| 2^121 (2.6e36).

| The backbone generator's period can be split up into a
| number of 'virtual generators' (A), each of which can be set
| to access a number of 'segments' (V) of length W, subject to
| the constraint that A * V * W <= 2^120.

| This is possible because the backbone generator has
| 'jump-ahead facilities' -- i.e. a quick method of calculating
| the state of the generator at a given point in its cycle
| without having to actually draw a large number of variates.

Algorithm | state(j) = seed(j)
for each | newstate(j) = (a(j) * oldstate(j)) mod m(j)
generator:  | output(j) = state(j)

Algorithm | output = (state(1) - state(2) + state(3) - state(4)) mod m(0)
for combination:|
```

```
Comments:  | For a discussion of PMMLCG generators, see the source file
| PMMLCGgen.h.

Configuration | unsigned int numGenerators; // # virtual generators
parameters:  | unsigned int numSegments; // log2(#segments/vGen)
| unsigned int segmentLength; // log2(segment length)

State:  | unsigned int state [numGenerators] [4];
| (each virtual generator has a 4-integer state vector)

Output types | unsigned int in [0,unsignedMax].
(for each | float in [0.0,1.0) using 1 iteration of the generator.
virtual | double in [0.0,1.0) using 1 or 2 iterations of the generator.
generator):  | long double in [0.0,1.0) using 2 iterations of the generator.

Output range:  | The combined generator yields output in the range [0,m0-2]
| (i.e. [0,2147483645].)

Limits on | Number of generators A <= 2^25  (33,554,432)
parameters:  | (You'll likely run out of memory before then ...)
| Number of segments / generator <= 2^63 (i.e. v<=63)
| Length of each segment <= 2^63 (i.e. w<=63)

Defaults:  | A=128, v=31, w=41

Valid seeds:  | Using create:setStateFromSeed: in [1, 2^32-1]
| Using create:setStateFromSeeds:
| The generator must be seeded with 4 starting seeds:
|  Seed[0]: [1, 2147483646] = [1,m1-1]
|  Seed[1]: [1, 2147483542] = [1,m2-1]
|  Seed[2]: [1, 2147483422] = [1,m3-1]
|  Seed[3]: [1, 2147483322] = [1,m4-1]

Cycles:  | With properly chosen parameters, the component generators
| have a single full cycle of length (m-1).
| The combined generator has a cycle length of
| 2147483646*2147483542*2147483422*2147483322/8 > 2^120.

Output Quality:  | All bits should be safe.
| The combined generator has much better statistical
| properties than either of its component PMMLCG's.

Reference for | L'Ecuyer and Andres, op. cit.
output quality:  |

Implementation  | The code for -getUnsignedMax uses a trick devised by
```

```
comment:        | Schrage to avoid 32-bit overflow when calculating the
| next state (i.e. using q=m/a and r=m%a as auxiliary
| values). Using 64-bit math (long long integers) is more
| straightforward. Unfortunately, on x86 at least, there is
| a significant speed penalty for doing so.
| The output, however, would be identical. YMMV.

C4LCGXgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       11.688 uS
|   -getFloatSample, getDoubleSample:         13.644 uS
|   -getFatDoubleSample, getLongDoubleSample: 26.154 uS

Relative speed:
| Speed 0.316 (time 3.160) relative to MT19937 getUnsignedSample
--------------- | -----------------------------------------------------------
---

Parameters:  | a[COMPONENTS] = {      45991,     207707,
|      138556,      49689 }
| m[COMPONENTS] = { 2147483647, 2147483543,
|  2147483423, 2147483323 }

Reference for | L'Ecuyer and Andres, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length close to 2^121 (2.6e36).

Output Quality:  | The authors performed an exhaustive search for component
| generators that would yield the best statistical quality.

Reference for | L'Ecuyer and Andres, op. cit.
output quality:  |
```

(./C3MWCgen.html)

(./Interface.html)

(./LCG1-3gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

LCG1-3gen Classes

```
//
// LCG1, LCG2, LCG3 --
//   single short random number generators,
//   NOT recommended for serious use; these are included for
//   historical reasons (compatibility with earlier releases).
//
@deftype LCGgen      <SingleShortGenerator> @end
@deftype LCG1gen        <LCGgen, CREATABLE> @end
@deftype LCG2gen        <LCGgen, CREATABLE> @end
@deftype LCG3gen        <LCGgen, CREATABLE> @end
```

Use of these generators is NOT Recommended

General Documentation

```
Generator name:  | LCG

Description:  | Linear Congruential Generator

Reference:  | This is a classic generator.
| See for example Donald Knuth,
| The Art of Computer Programming,
| vol. II, "Seminumerical Algorithms."

| state = seed
Algorithm:  | newstate = ((a * oldstate) + c ) mod 2^32
| output = state

Comments:  | This generator relies on controlled overflow at 32 bits.
| This requires that unsigned be a 32bit value that follows
| ANSI C rules.
| Knuth claims that the adder c does not matter much, as
| long as it has no factors in common with the modulus 2^32.


Parameters:  | unsigned int c // adder
| unsigned int a // multiplier
```

```
State:  | unsigned int state // a "state vector" of length 1

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,2^32-1] may be returned
| (i.e. [0,4294967295].)

Valid seeds:  | Using create:setStateFromSeed:  [1,2^32-1]
| Using create:setStateFromSeeds: [1,2^32-1]

Cycles:  | With properly chosen parameters, these generators
| have a single full cycle of length (2^32).

Output Quality:  | Low order bits are poor.

Reference for | Knuth, op.cit.
output quality:  |

LCGgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       2.564 uS
|   -getFloatSample, getDoubleSample:         4.295 uS
|   -getFatDoubleSample, getLongDoubleSample: 7.734 uS

Relative speed:
| Speed 1.442 (time 0.693) relative to MT19937 getUnsignedSample
```

LCG1gen Documentation

```
Parameters:  | a =     1,664,525
| c = 1,013,904,223

Reference for | The multiplier is from Knuth vol II (line 26 of table).
parameters:  | Adder is from Numerical Recipes, in turn from H.W.Lewis

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length m.

Output Quality:  | Low-order bits are poor.

Reference for | Knuth, op.cit. and Numerical Recipes.
output quality:  |
```

LCG2gen Documentation

```
Parameters:  | a =        69,069
| c = 1,013,904,223

Reference for | The multiplier is from Knuth vol II (line 25 of table).
parameters:  | Adder is from Numerical Recipes, in turn from H.W.Lewis

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length m.

Output Quality:  | Low-order bits are poor.

Reference for | Knuth, op.cit. and Numerical Recipes.
output quality:  |
```

LCG3gen Documentation

```
Parameters:  | a =   1,664,525
| c = 152,193,325

Reference for | The multiplier is from Knuth vol II (line 26 of table).
parameters:  | Adder is from Numerical Recipes, in turn from H.W.Lewis

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length m.

Output Quality:  | Low-order bits are poor.

Reference for | Knuth, op.cit. and Numerical Recipes.
output quality:  |
```

(./C4LCGXgen.html)

(./Interface.html)

(./MRG5-7gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

MRG5-7gen Classes

```
//
// MRG5, MRG6, MRG7 --
//   single long generators recommended for use.
//
@deftype MRGgen      <SingleLongGenerator>  @end
@deftype MRG5gen      <MRGgen, CREATABLE> @end
@deftype MRG6gen      <MRGgen, CREATABLE> @end
@deftype MRG7gen      <MRGgen, CREATABLE> @end
```

General Documentation

```
Generator name:  | MRG

Description:  | Multiple Recursive LCG Generator

Reference:  | Pierre L'Ecuyer, Francois Blouin and Raymond Couture,
| "A Search for Good Multiple Recursive Random Generators."
| See <http://www.iro.umontreal.ca/~lecuyer>.

Algorithm:  | x_n = (a_1*x_n-1 + a_2*x_n-2 ... + a_k*x_n-k) mod m

Comments:  | These generators require k multipliers and k past values
| to be kept. In their paper, the authors investigate MRG's
| of order k from 1 to 7. They provide several sets of
| parameters which they recommend out of a large number that
| were tested. Generally, the quality of the generators
| increases with k.

Implementation:  | The authors discovered that most of the benefit of these
| generators can be had if only a_1 and a_k are nonzero and the
| rest zero; this yields faster execution since only 2 terms
| need be calculated. We implement 4 such generators here,
| all with modulus m close to 2^31.


Parameters:  | unsigned int m // modulus
| unsigned int a1, ak // multipliers
| unsigned int k // the order of the generator

State:  | unsigned int x[k] // state vector
```

```
Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,m-2].

Valid seeds:  | Using create:setStateFromSeed:  [1, 2^32-1]
| Using create:setStateFromSeeds: [1, unsignedMax]
| (unsignedMax differs between MRG5, MRG6, MRG7)

Cycles:  | With properly chosen parameters, these generators
| have a single cycle of length m^k - 1.
| Thus, for m ~= 2^31 we have:
| k=5 => cycle ~= 2^155,
| k=6 => cycle ~= 2^186,
| k=7 => cycle ~= 2^217.

Output Quality:  | Good.

Reference for | L'Ecuyer et al, op. cit.
output quality:  |

MRG5gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        9.674 uS
|   -getFloatSample, getDoubleSample:         11.524 uS
|   -getFatDoubleSample, getLongDoubleSample: 22.135 uS

Relative speed:
| Speed 0.382 (time 2.616) relative to MT19937 getUnsignedSample
-------------- . -----------------------------------------------------------
--

MRG6gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       10.449 uS
|   -getFloatSample, getDoubleSample:         12.337 uS
|   -getFatDoubleSample, getLongDoubleSample: 27.735 uS

Relative speed:
| Speed 0.354 (time 2.826) relative to MT19937 getUnsignedSample
-------------- . -----------------------------------------------------------
--

MRG7gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       10.913 uS
```

```
|   -getFloatSample, getDoubleSample:         12.764 uS
|   -getFatDoubleSample, getLongDoubleSample: 24.627 uS

Relative speed:
| Speed 0.339 (time 2.951) relative to MT19937 getUnsignedSample
```

MRG5gen Documentation

```
Parameters:  | unsigned int k = 5 // lag order
| unsigned int m = 2^31 - 1 // modulus
| unsigned int a1 = 107374182 // multiplier
| unsigned int a5 = 104480 // multiplier

Reference for | L'Ecuyer et al, op. cit., Table 3.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (2^31-1)^5 - 1,
| i.e. 2^154 < cycle < 2^155.

Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | L'Ecuyer et al, op. cit.
output quality:  |
```

MRG6gen Documentation

```
Parameters:  | unsigned int k = 6 // lag order
| unsigned int m = 2^31 - 1 // modulus
| unsigned int a1 = 177786 // multiplier
| unsigned int a6 = 64654 // multiplier

Reference for | L'Ecuyer et al, op. cit., Table 3.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (2^31-1)^6 - 1,
| i.e. 2^185 < cycle < 2^186.

Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | L'Ecuyer et al, op. cit.
output quality:  |
```

MRG7gen Documentation

```
Parameters:  | unsigned int k = 7 // lag order
| unsigned int m = 2^31 - 19 // modulus
| unsigned int a1 = 1071064 // multiplier
| unsigned int a7 = 2113664 // multiplier

Reference for | L'Ecuyer et al, op. cit., Table 3.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (2^31-1)^7 - 1,
| i.e. 2^216 < cycle < 2^217.

Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | L'Ecuyer et al, op. cit.
output quality:  |
```

(./LCG1-3gen.html)

(./Interface.html)

(./MT19937gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

MT19937gen Class

```
//
// MT19937 --
//   single *very* long generator recommended for use.
//
@deftype MT19937gen  <SingleLongGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | MT19937

Description:  | Twisted GFSR generator

Reference:  | Makoto Matsumoto and ??? Nishimura,
| code posted to the pLab web site, xx/xx/97.
| Site: <http://random.mat.sbg.ac.at>
| Code: <ftp://random.mat.sbg.ac.at/pub/data/mt19937b-int.c>
| Article reportedly submitted to ACM TOMACS.

|
Algorithm:  |
|

Comments:  |

Implementation:  |

Parameters:  | unsigned int N // size of state array
| unsigned int M // lag
| unsigned int a // "Magic numbers"
| unsigned int b,c //    "     "
| unsigned int u,s,t,l //    "     "

State:  | unsigned int state[N] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)
```

```
Output range:  | Any number in the range [0,2^32] may be returned

Valid seeds:  | Using create:setStateFromSeed:  in [1,2^32-1]
| Using create:setStateFromSeeds: in [1,2^32-1].

Cycles:  | This generator has a single cycle of length 2^19937-1.

Output Quality:  | Excellent equidistribution properties in all dimensions
| up to 623 are claimed.

Reference for | Matsumoto and Nishimura, op. cit.
output quality:  | See also pLab Report #07 on the pLab web site.

MT19937gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        3.698 uS
|   -getFloatSample, getDoubleSample:          5.502 uS
|   -getFatDoubleSample, getLongDoubleSample: 10.159 uS

Relative speed:
| Speed 1.0 (time 1.0) relative to itself (it's the benchmark)

--------------- . -----------------------------------------------------------
--

Parameters:  | unsigned int w=32 // word length
| unsigned int N=624 // size of state array
| unsigned int M=397 // lag
| unsigned int s=7 // 'magic vector'
| unsigned int t=15 //    "     "
| unsigned int u=11 //    "     "
| unsigned int l=18 //    "     "
| unsigned int b=0x102d1200 //    "     "
| unsigned int c=0x66e50000 //    "     "
| unsigned int a=0x6b5eccf6 //    "     "

Reference for | Matsumoto and Nishimura, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length 2^19937 -1,
| or 10^6001 !!! (The longest period yet.)

Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | Matsumoto and Nishimura, op. cit.
output quality:  |
```

(./MRG5-7gen.html)

(./Interface.html)

(./MWCA-Bgen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

MWCA-Bgen Classes

```
//
// MWC --
//   two long generators recommended for use.
//
@deftype MWCAgen    <SingleLongGenerator, CREATABLE> @end
@deftype MWCBgen    <SingleLongGenerator, CREATABLE> @end
```

MWCAgen Documentation

```
Generator name:  | MWCA

Description:  | Multiply-With-Carry generator

Reference:  | Article posted to Usenet in 1994 by George Marsaglia,
| "The Mother of all Random Number Generators."
| Source found at <ftp://ftp.taygeta.com/pub/c/mother.c>
| Also document /papers/mwc1.ps on his Diehard CD-ROM,
| found at <http://www.hku.hk/internet/randomCD.html>.

Algorithm:  | K = a * (K & 65535) + (K>>16);
| J = b * (J & 65535) + (J>>16);
| return (K<<16) + J;

Comments:  | This generator is a combination of 2 16-bit
| Multiply-With-Carry generators.

Implementation:  |

Parameters:  | unsigned int a = 30903 // multiplier
| unsigned int b = 18000 // multiplier

State:  | unsigned int K,J // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,2^32-1].
```

```
Valid seeds:  | Using create:setStateFromSeed:  [1, 2^32-1]
| Using create:setStateFromSeeds: each component may have
| a 'state' <= 65535 and a 'carry' <= (multiplier-1). Thus
| the first component has a maximum seed of
| ( (a-1) << 16) + 65535.

Cycles:  | This generator is claimed to be strictly periodic, with a
| period > 2^59. (There's possibly two such cycles.)

Output Quality:  | Good.

Reference for | Marsaglia, op. cit.
output quality:  |

MWCgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       3.399 uS
|   -getFloatSample, getDoubleSample:         5.131 uS
|   -getFatDoubleSample, getLongDoubleSample: 9.447 uS

Relative speed:
| Speed 1.088 (time 0.919) relative to MT19937 getUnsignedSample
--------------- . ------------------------------------------------------------
--
```

MWCBgen Documentation

```
Generator name:  | MWCB

Description:  | Multiply-With-Carry generator

Reference:  | Article posted to Usenet in 1994 by George Marsaglia,
| "The Mother of all Random Number Generators."
| Source found at <ftp://ftp.taygeta.com/pub/c/mother.c>
| Also document /papers/mwc1.ps on his Diehard CD-ROM,
| found at <http://www.hku.hk/internet/randomCD.html>.

Algorithm:  | K = a * (K & 65535) + (K>>16);
| J = b * (J & 65535) + (J>>16);
| return (K<<16) + (J&65535);

Comments:  | This generator is a combination of 2 16-bit
| Multiply-With-Carry generators.

Implementation:  | This generator implements an alternate manner of
| conjoining the two components (differs from MWCA).
```

| It also passes all tests and is fast.

Parameters: | unsigned int a = 30903 // multiplier
| unsigned int b = 18000 // multiplier

State: | unsigned int K,J // state vector

Output types: | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range: | Any number in the range [0,unsignedMax] = [0,2^32-1].

Valid seeds: | Using create:setStateFromSeed: [1, 2^32-1]
| Using create:setStateFromSeeds: each component may have
| a 'state' <= 65535 and a 'carry' <= (multiplier-1). Thus
| the first component has a maximum seed of
| ((a-1) << 16) + 65535.

Cycles: | This generator is claimed to be strictly periodic, with a
| period > 2^59. (There's possibly two such cycles.)

Output Quality: | Good.

Reference for | Marsaglia, op. cit.
output quality: |

MWCgen | On a 486/66, the following approximate measures were obtained:
speed: | -getUnsignedSample: 3.399 uS
| -getFloatSample, getDoubleSample: 5.131 uS
| -getFatDoubleSample, getLongDoubleSample: 9.447 uS

Relative speed:
| Speed 1.088 (time 0.919) relative to MT19937 getUnsignedSample

---

(./MT19937gen.html)

(./Interface.html)

(./PMMLCG1-9gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

PMMLCG 1-9 Classes

```
//
// PMMLCG1, PMMLCG2, PMMLCG3 --
//   single short random number generators recommended for use.
//
// PMMLCG4, PMMLCG5, PMMLCG6, PMMLCG7 --
//   component generators of C4LCGX
//
// PMMLCG8, PMMLCG9 --
//   component generators of C2LCGX
//
@deftype PMMLCGgen   <SingleShortGenerator> @end
@deftype PMMLCG1gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG2gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG3gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG4gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG5gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG6gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG7gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG8gen  <PMMLCGgen, CREATABLE> @end
@deftype PMMLCG9gen  <PMMLCGgen, CREATABLE> @end
```

General Documentation

```
Generator name:  | PMMLCG

Description:  | Prime Modulus Multiplicative Linear Congruential Generator

Reference:  | Stephen K. Park and Keith W. Miller
| (CACM 31:10, October 1988, 1192-1201),
| "Random Number Generators: Good Ones Are Hard To Find."
| The authors refer to this generator as a "minimally
| acceptable" or "minimally sufficient" generator.

| state = seed
Algorithm:  | newstate = (a * oldstate) mod m
| output = state-1

Comments:  | The modulus m must be a prime for this class of generator
| (as opposed to the more common power of two modulus).
| For these generators we use m=2^31-1 (2,147,483,647).
```

```
| The choice of a good multiplier a is crucial.

Parameters:  | unsigned int m // modulus
| unsigned int a // multiplier

State:  | unsigned int state // a "state vector" of length 1

Output types:  | unsigned int (we return the current state-1).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | The basic algorithm yields output in the range [1,m-1].
| We subtract 1 to return output in the range [0,m-2]
| (i.e. [0,2147483645].)

Valid seeds:  | create:setStateFromSeed:  [1, 2^32-1]
| create:setStateFromSeeds: [1, m-1]

Cycles:  | With properly chosen parameters, these generators
| have a single full cycle of length (m-1).
| (State=0 is an absorbing state which will not normally
| occur; we avoid setting the starting state to 0.)

Output Quality:  | All bits should be safe.
| (I.e. low order bits should not have small cycles.)
| With m prime, this generator is better than standard LCG.

Reference for | Park and Miller, op. cit.
output quality:  |

Implementation  | The code for -getUnsignedMax uses a trick devised by
comment:        | Schrage to avoid 32-bit overflow when calculating the
| next state (i.e. using q=m/a and r=m%a as auxiliary
| values). Using 64-bit math (long long integers) is more
| straightforward. Unfortunagely, on x86 at least, there
| is a 35% speed penalty for doing so. The output, however,
| is identical. YMMV.

PMMLCGgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        4.715 uS
|   -getFloatSample, getDoubleSample:          6.500 uS
|   -getFatDoubleSample, getLongDoubleSample: 12.152 uS

Relative speed:
| Speed 0.792 (time 1.263) relative to MT19937 getUnsignedSample
```

--------------- | --

PMMLCG1gen Documentation

```
Parameters:  | a =        16,807
| m = 2,147,483,647 (2^31-1)

Reference for | Park and Miller, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.
| These parameters have been extensively tested by the authors.

Reference for | Park and Miller, op. cit.
output quality:  |
```

PMMLCG2gen Documentation

```
Parameters:  | a =        48,271
| m = 2,147,483,647 (2^31-1)

Reference for | Park and Miller, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.
| These parameters are not as well tested as PMMLCG1,
| but there's reason to believe they may be better.

Reference for | Park and Miller, op. cit.
output quality:  |
```

PMMLCG3gen Documentation

```
Parameters:  | a =        69,621
| m = 2,147,483,647 (2^31-1)

Reference for | Park and Miller, op. cit.
```

```
choice of  |
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.
| These parameters are not as well tested as PMMLCG1,
| but there's reason to believe they may be better.

Reference for | Park and Miller, op. cit.
output quality:  |
```

PMMLCG4gen Documentation

```
Parameters:  | a =        45,991
| m = 2,147,483,647 (2^31-1)

Reference for | Pierre L'Ecuyer and Terry H. Andres,
choice of  | "A Random Number Generator Based on the Combination
parameters:  | of Four LCGs", unpubl.

| This is one of the component generators of their CLCG4.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Andres, op.cit.
output quality:  |
```

PMMLCG5gen Documentation

```
Parameters:  | a =       207,707
| m = 2,147,483,543 (2^31-105)

Reference for | Pierre L'Ecuyer and Terry H. Andres,
choice of  | "A Random Number Generator Based on the Combination
parameters:  | of Four LCGs", unpubl.

| This is one of the component generators of their CLCG4.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).
```

```
Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Andres, op.cit.
output quality:  |
```

PMMLCG6gen Documentation

```
Parameters:  | a =       138,556
| m = 2,147,483,423 (2^31-225)

Reference for | Pierre L'Ecuyer and Terry H. Andres,
choice of  | "A Random Number Generator Based on the Combination
parameters:  | of Four LCGs", unpubl.

| This is one of the component generators of their CLCG4.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Andres, op.cit.
output quality:  |
```

PMMLCG7gen Documentation

```
Parameters:  | a =        49,689
| m = 2,147,483,323 (2^31-325)

Reference for | Pierre L'Ecuyer and Terry H. Andres,
choice of  | "A Random Number Generator Based on the Combination
parameters:  | of Four LCGs", unpubl.

| This is one of the component generators of their CLCG4.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Andres, op.cit.
output quality:  |
```

PMMLCG8gen Documentation

```
Parameters:  | a =        40,014
| m = 2,147,483,563 (2^31-85)

Reference for | Pierre L'Ecuyer and Serge Cote,
choice of  | "Implementing a Random Number Package with Splitting
parameters:  | Facilities," ACM Trans. Math. Softw. 17 no 1 (March 1991).

| This is one of the component generators of their generator.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Cote, op.cit.
output quality:  |
```

PMMLCG9gen Documentation

```
Parameters:  | a =        40,692
| m = 2,147,483,399 (2^31-249)

Reference for | Pierre L'Ecuyer and Serge Cote,
choice of  | "Implementing a Random Number Package with Splitting
parameters:  | Facilities," ACM Trans. Math. Softw. 17 no 1 (March 1991).

| This is one of the component generators of their generator.

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length (m-1).

Output Quality:  | All bits should be safe.

Reference for | L'Ecuyer and Cote, op.cit.
output quality:  |
```

(./MWCA-Bgen.html)

(./Interface.html)

(./PSWBgen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

PSWBgen Class

```
//
// PSWB --
//   single long generator recommended for use.
//
@deftype PSWBgen   <SingleLongGenerator, CREATABLE> @end
```

Documentation

```
Generator name:  | PSWB

Description:  | Subtract-with-Borrow Congruential Generator
|    with Prime Modulus

Reference:  | George Marsaglia and Arif Zaman,
| "A New Class of Random Number Generators",
| Annals of Applied Probability vol 1 no 3 (1991), 462-480.
| SWB is an improvement on the Lagged Fibonacci generators.

| PSWB is an improvement on SWB in that the use of a prime
| modulus guarantees a single full cycle.
| It's slower, of course.

| state vector = f(seed)
Algorithm:  | state_n = (state_(n-s) - state_(n-r) - carry) mod m; r > s
| output = state_n

| These generators use a basic algorithm of the form
| X_n = f(X_(n-r),X_(n-s)) mod m; r>s
| The function f is typically xor, addition, subtraction,
| multiplication or subtraction with carry. It uses
| simpler math than a basic LCG, but keeps a larger state.

Implementation:  | The state vector contains the generator's r most recent
| outputs (X_n .. X_(n-r)). This vector is initialized from
| an LCG generator using one 32-bit seed value. After
| initialization we throw away r*5 numbers to 'prime' the
| generator (i.e. make sure we have entered a closed cycle.)

Parameters:  | unsigned int r // long lag
| unsigned int s // short lag
```

```
| unsigned int m // modulus

State:  | unsigned int state[r] // state vector
| int carry // carry bit

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,m-1] may be returned.

Valid seeds:  | Using create:setStateFromSeed:  [1,2^32-1]
| Using create:setStateFromSeeds: r seeds in [1,2^32-7]
| and one seed in [0,1] (the carry)

Cycles:  | With properly chosen parameters, this generator
| has a single cycle of length close to 10^414.

| Unfortunately, there are no 'jump-ahead' or 'splitting'
| facilities that would allow us to jump ahead in the cycle,
| other than calling the generator a number of times.

Output Quality:
| Much better than other Lagged Fibonacci generators, such as
| ACG and SCG.

Reference for | Marsaglia and Zaman, op.cit.
output quality:  |

PSWBgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                      3.452 uS
|   -getFloatSample, getDoubleSample:         5.166 uS
|   -getFatDoubleSample, getLongDoubleSample: 9.544 uS

Relative speed:
| Speed 1.070 (time 0.933) relative to MT19937 getUnsignedSample
--------------- | -----------------------------------------------------------
---

Parameters:  | r = 43 // long lag
| s = 22 // short lag
| m = 2^32 - 5 // modulus (prime)

Reference for | Marsaglia, op. cit.
parameters:  |
```

Cycles: | With the chosen parameters, this generator
| has a single cycle of length close to 10^414.

Output Quality: | Good.

Reference for | Marsaglia, op.cit., p. 479
output quality: |

(./PMMLCG1-9gen.html)

(./Interface.html)

(./RWC28gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

RWC2,8gen Classes

```
//
// RWC2 --
// RWC8 ("Mother") --
//   single long generator recommended for use.
//
@deftype RWC2gen   <SingleLongGenerator, CREATABLE> @end
@deftype RWC8gen   <CombinedLongGenerator, CREATABLE> @end
```

RWC2gen Documentation

```
Generator name:  | RWC2

Description:  | 2-lag Recursion-With-Carry generator

Reference:  | Article posted to Usenet in 1994 by George Marsaglia,
| "The Mother of all Random Number Generators."
| Source found at <ftp://ftp.taygeta.com/pub/c/mother.c>
| Also document /papers/rwc1.ps on his Diehard CD-ROM,
| found at <http://www.hku.hk/internet/randomCD.html>.

Algorithm:  | x[n]  = a*(x[n-1]+x[n-2]) + carry % 2^32
| carry = a*(x[n-1]+x[n-2]) + carry / 2^32

Comments:  | This generator is a 2-lag MWC generator

Implementation:  | This generator is implemented using 64-bit math

Parameters:  | unsigned int a = 1,111,111,464 // multiplier

State:  | unsigned int carry,x1,x2 // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,2^32-1].

Valid seeds:  | Using create:setStateFromSeed:  [1, 2^32-1]
```

```
| Using create:setStateFromSeeds: x1 and x2: [1, 2^32-1]
| The carry can only be [1,a-1].

Cycles:  | This generator is claimed to have a period > 2^92.
| It may have rho sequences up to length r=2.
| It may have multiple cycles of the stated period length.

Output Quality:  | Good.

Reference for | Marsaglia, op. cit.
output quality:  |

RWC2gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                        7.445 uS
|   -getFloatSample, getDoubleSample:          8.845 uS
|   -getFatDoubleSample, getLongDoubleSample: 17.338 uS

Relative speed:
| Speed 0.497 (time 2.013) relative to MT19937 getUnsignedSample
```

RWC8gen Documentation

```
Generator name:  | RWC8

Description:  | 8-lag Recursion-With-Carry generator
| ("The Mother of all RNG's")

Reference:  | Article posted to Usenet in 1994 by George Marsaglia,
| "The Mother of all Random Number Generators."
| Source found at <ftp://ftp.taygeta.com/pub/c/mother.c>
| Also document /papers/mwc1.ps on his Diehard CD-ROM,
| found at <http://www.hku.hk/internet/randomCD.html>.

Algorithm:  | x = sum(ai*xi) + carry mod 2^16
| y = sum(bi*yi) + carry mod 2^16
| return ((x<<16)+y)

Comments:  | This generator is a combination of 2 16-bit
| 8-lag Recursion-With-Carry generators.

Parameters:  | unsigned int a[8] =
| { 1941, 1860, 1812, 1776, 1492, 1215, 1066, 12013 }

| unsigned int b[8] =
| { 1111, 2222, 3333, 4444, 5555, 6666, 7777,  9272 }
```

```
State:  | unsigned short m1[10], m2[10] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,unsignedMax] = [0,2^32-1].

Valid seeds:  | Using create:setStateFromSeed:  [1, 2^32-1]
| Using create:setStateFromSeeds:
|   a total of 18 16-bit seeds are needed;
|   each xi may be in [1, 2^31-1].
|   each carry may be in [1, sum(multipliers)-1].

Cycles:  | This generator is claimed to have a period > 2^250.
| It may have rho sequences up to length r = 8.
| It may have multiple cycles of the stated period length.

Output Quality:  | Claimed to be good.

Reference for | Marsaglia, op. cit.
output quality:  |

RWC8gen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       14.649 uS
|   -getFloatSample, getDoubleSample:         16.400 uS
|   -getFatDoubleSample, getLongDoubleSample: 31.915 uS

Relative speed:
| Speed 0.252 (time 3.961) relative to MT19937 getUnsignedSample
--------------- . ------------------------------------------------------------
--
```

(./PSWBgen.html)

(./Interface.html)

(./SCGgen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

SCGgen Class

```
//
// SCG --
//   single long random number generator,
//   NOT recommended for serious use; included for
//   historical reasons (compatibility with earlier releases).
//
@deftype SCGgen     <SingleLongGenerator, CREATABLE> @end
```

Use of this generator is NOT Recommended

Documentation

```
Generator name:  | SCG

Description:  | Subtractive Congruential Generator

Reference:  | SCG is an often used algorithm, mentioned in Knuth,
| Numerical Recipes, etc.

| state vector = f(seed)
Algorithm:  | state_n = (state_(n-r) - state_(n-s)) mod 10^9; r > s
| output = state_n

| SCG is in the Lagged Fibonacci class of generators.
| These generators use a basic algorithm of the form
| X_n = f(X_(n-r),X_(n-s)) mod m; r>s
| The function f is typically xor, addition, subtraction,
| multiplication or subtraction with carry. It uses
| simpler math than a basic LCG, but keeps a larger state.

Implementation:  | The state vector contains the generator's r most recent
| outputs (X_n .. X_(n-r)). This vector is initialized from
| an LCG generator using one 32-bit seed value. After
| initialization we throw away r*5 numbers to 'prime' the
| generator (i.e. make sure we have entered a closed cycle.)

Parameters:  | unsigned int r // long lag
```

```
| unsigned int s // short lag
| unsigned int m // modulus

State:  | unsigned int state[r] // state vector

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,10^9-1] may be returned
| (i.e. [0,999999999].)

Valid seeds:  | Using create:setStateFromSeed:  [1,2^32-1]
| Using create:setStateFromSeeds: [1,999999998]

Cycles:  | With properly chosen parameters, these generators
| have a cycle of length about 2^r.

Output Quality:  | Low order bits are poor, and there are other problems.
| SWB is a better generator in the same class, and it is
| only slightly slower.

Reference for | Knuth, op.cit.
output quality:  |

SCGgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       3.328 uS
|   -getFloatSample, getDoubleSample:         5.058 uS
|   -getFatDoubleSample, getLongDoubleSample: 9.274 uS

Relative speed:
| Speed 1.111 (time 0.900) relative to MT19937 getUnsignedSample
--------------- | -----------------------------------------------------------
---

Parameters:  | r = 55
| s = 24
| m = 10^9

Reference for |
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length ~ 2^r.
```

Output Quality:	Low-order bits are poor.
Reference for output quality:	Knuth, op.cit.

(./RWC28gen.html)

(./Interface.html)

(./SWB1-3gen.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

SWB1-3gen Classes

```
//
// SWB1, SWB2, SWB3 --
//   single long generators recommended for use.
//
@deftype SWBgen   <SingleLongGenerator> @end
@deftype SWB1gen    <SWBgen, CREATABLE> @end
@deftype SWB2gen    <SWBgen, CREATABLE> @end
@deftype SWB3gen    <SWBgen, CREATABLE> @end
```

General Documentation

```
Generator name:  | SWB

Description:  | Subtract-with-Borrow Congruential Generator

Reference:  | George Marsaglia and Arif Zaman,
| "A New Class of Random Number Generators",
| Annals of Applied Probability vol 1 no 3 (1991), 462-480.
| SWB is an improvement on the Lagged Fibonacci generators.

| state vector = f(seed)
Algorithm:  | state_n = (state_(n-s) - state_(n-r) - carry) mod 2^32; r > s
| output = state_n

| These generators use a basic algorithm of the form
| X_n = f(X_(n-r),X_(n-s)) mod m; r>s
| The function f is typically xor, addition, subtraction,
| multiplication or subtraction with carry. It uses
| simpler math than a basic LCG, but keeps a larger state.

Note:  | This generator relies on controlled overflow at 32 bits.
| This requires that unsigned be a 32bit value that follows
| ANSI C rules.

Implementation:  | The state vector contains the generator's r most recent
| outputs (X_n .. X_(n-r)). This vector is initialized from
| an LCG generator using one 32-bit seed value. After
| initialization we throw away r*5 numbers to 'prime' the
| generator (i.e. make sure we have entered a closed cycle.)
```

```
Parameters:  | unsigned int r // long lag
| unsigned int s // short lag

State:  | unsigned int state[r] // state vector
| int carry // carry bit

Output types:  | unsigned int (we return the current state).
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,2^32-1] may be returned
| (i.e. [0,4294967295].)

Valid seeds:  | Using create:setStateFromSeed:  [1,2^32-1]
| Using create:setStateFromSeeds: r seeds in [1,2^32-2]
| and one seed in [0,1] (the carry)

Cycles:  | With properly chosen parameters, these generators
| have a number of long cycles, up to 10^354 in length.
| UNFORTUNATELY there is no way to pick different
| cycles, nor any easy way to 'jump ahead' in the
| chosen cycle (other than calling -getUnsignedSample
| a number of times).

Output Quality:
| Much better than other Lagged Fibonacci generators, such as
| ACG and SCG.

Reference for | Marsaglia and Zaman, op.cit.
output quality:  |

SWBgen | On a 486/66, the following approximate measures were obtained:
speed:  |   -getUnsignedSample:                       3.285 uS
|   -getFloatSample, getDoubleSample:         5.057 uS
|   -getFatDoubleSample, getLongDoubleSample: 9.209 uS

Relative speed:
| Speed 1.126 (time 0.888) relative to MT19937 getUnsignedSample
-------------- | -----------------------------------------------------------
---
```

SWB1gen Documentation

```
Parameters:  | r = 37
```

```
| s = 24

Reference for | Marsaglia, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has 64 cycles of length 10^354

Output Quality:  | Good.

Reference for | Marsaglia, op.cit., p. 479
output quality:  |
```

SWB2gen Documentation

```
Parameters:  | r = 24
| s = 19

Reference for | Marsaglia, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has 1536 cycles of length 10^228

Output Quality:  | Good.

Reference for | Marsaglia, op.cit., p. 479
output quality:  |
```

SWB3gen Documentation

```
Parameters:  | r = 21
| s =  6

Reference for | Marsaglia, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has 192 cycles of length 10^200

Output Quality:  | Good.

Reference for | Marsaglia, op.cit., p. 479
output quality:  |
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

TGFSRgen Classes

```
//
// TT403, TT775, TT800 --
//   single long generators recommended for use.
//
@deftype TGFSRgen    <SingleLongGenerator>  @end
@deftype TT403gen     <TGFSRgen, CREATABLE> @end
@deftype TT775gen     <TGFSRgen, CREATABLE> @end
@deftype TT800gen     <TGFSRgen, CREATABLE> @end
```

General Documentation

```
Generator name:  | TGFSR

Description:  | Twisted GFSR generator

Reference:  | Makoto Matsumoto and Yoshiharu Kurita,
| "Twisted GFSR Generators II".
| ACM TOMACS, vol 4 no 3 July 1994 pp. 254-266.
| Amended by M. Matsumoto, July 8 1996.

|
Algorithm:  |
|

Comments:  |

Implementation:  |

Parameters:  | unsigned int w // word length
| unsigned int N // size of state array
| unsigned int M // lag
| unsigned int s // 'magic vector'
| unsigned int t //    "     "
| unsigned int b //    "     "
| unsigned int c //    "     "
| unsigned int a //    "     "

State:  | unsigned int state[N] // state vector

Output types:  | unsigned int (we return the current state).
```

```
| float in [0.0,1.0) using 1 iteration of the generator.
| double in [0.0,1.0) using 1 or 2 iterations of the generator.
| long double in [0.0,1.0) using 2 iterations of the generator.
| (See the implementation file for more details.)

Output range:  | Any number in the range [0,2^w-1] may be returned

Valid seeds:  | Using create:setStateFromSeed:  in [1,2^32-1]
| Using create:setStateFromSeeds: in [1,2^w-1].

Cycles:  | With properly chosen parameters, these generators
| have a single cycle of length 2^(w*N) -1.

Output Quality:  | Good.

Reference for | Matsumoto and Kurita, op. cit.
output quality:  | Also recommended by P. L'Ecuyer (private communication).

TT403gen, | On a 486/66, the following approximate measures were obtained:
TT775gen, |   -getUnsignedSample:                          4.654 uS
TT800gen |   -getFloatSample, getDoubleSample:          6.430 uS
speed:  |   -getFatDoubleSample, getLongDoubleSample: 11.953 uS

Relative speed:
| Speed 0.795 (time 1.259) relative to MT19937 getUnsignedSample
-------------- . -----------------------------------------------------------
--
```

TT403gen Documentation

```
Parameters:  | unsigned int w=31 // word length
| unsigned int N=13 // size of state array
| unsigned int M=2 // lag
| unsigned int s=8 // 'magic vector'
| unsigned int t=14 //    "     "
| unsigned int b=0x102d1200 //    "     "
| unsigned int c=0x66e50000 //    "     "
| unsigned int a=0x6b5eccf6 //    "     "

Reference for | Matsumoto and Kurita, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length 2^403 -1.

Output Quality:  | Good. Parameters chosen on the basis of theory
```

```
| and empirical testing.

Reference for | Matsumoto and Kurita, op. cit.
output quality:  |
```

TT775gen Documentation

```
Parameters:  | unsigned int w=31 // word length
| unsigned int N=25 // size of state array
| unsigned int M=8 // lag
| unsigned int s=6 // 'magic vector'
| unsigned int t=14 //    "     "
| unsigned int b=0x1abd5900 //    "     "
| unsigned int c=0x776a0000 //    "     "
| unsigned int a=0x6c6cb38c //    "     "

Reference for | Matsumoto and Kurita, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length 2^775 -1.

Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | Matsumoto and Kurita, op. cit.
output quality:  |
```

TT800gen Documentation

```
Parameters:  | unsigned int w=32 // word length
| unsigned int N=25 // size of state array
| unsigned int M=7 // lag
| unsigned int s=7 // 'magic vector'
| unsigned int t=15 //    "     "
| unsigned int b=0x2b5b2500 //    "     "
| unsigned int c=0xdb8b0000 //    "     "
| unsigned int a=0x8ebfd028 //    "     "

Reference for | Matsumoto and Kurita, op. cit.
parameters:  |

Cycles:  | With the chosen parameters, this generator
| has a single full cycle of length 2^800 -1.
```

```
Output Quality:  | Good. Parameters chosen on the basis of theory
| and empirical testing.

Reference for | Matsumoto and Kurita, op. cit.
output quality:  |
```

(./SWB1-3gen.html)

(./Interface.html)

(./Prot.html#ProbabilityDist)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

RandomBitDistribution Class

```
//
// RandomBitDistribution --
//   A generator that returns uniformly distributed single bit values
//   (i.e. fair coin tosses).
//
@deftype RandomBitDist <BooleanDistribution, CREATABLE>
USING
- (BOOL) getCoinToss;
@end
```

Documentation

```
Name:  | RandomBitDistribution

Description:  | Distribution returning YES and NO with equal probability

Algorithm:  | Extract a middle bit from an unsigned random number

Reference:  | (none)

Parameters:  | (none)

State:  | unsigned long long int currentCount

RandomBitDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getBooleanSample:                         5.656 uS
|  -average # of generator calls per variate: 1.000 (U)
|  -distribution net time:                    1.958 uS

Relative speed:
| Speed 0.654 (time 1.529) relative to MT19937 getUnsignedSample
--------------- . -----------------------------------------------------------
--
```

(Interface.html)

(BernoulliDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

BernoulliDistribution Class

```
//
// BernoulliDistribution --
distribution returning YES with given probability p
//
@deftype BernoulliDist <BooleanDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
                setProbability: (double) p;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setProbability: (double) p;
USING
- setProbability: (double) p;
- (double) getProbability;
- (BOOL)        getSampleWithProbability: (double) p;
@end
```

Documentation

```
Name:  | BernoulliDistribution

Description:  | Distribution returning binary value YES with probability p

Algorithm:  | Rejection method

Reference:  | Code contributed by Barry McMullin

Parameters:  | double Probability

State:  | unsigned long long int currentCount

BernoulliDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getBooleanSample:                        8.156 uS
|  -average # of generator calls per variate: 1.000 (tD)
|  -distribution net time:                   2.654 uS

Relative speed:
| Speed 0.453 (time 2.206) relative to MT19937 getUnsignedSample
```

--------------- . --
--

(./RandomBitDist.html)

(./Interface.html)

(./UniformIntegerDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

UniformIntegerDist Class

```
//
// UniformInteger --
//   A generator of integral values uniformly distributed across a closed
//   interval [min,max]. (The interval includes both its endpoints.)
//   Setting minValue == maxValue is allowed (and returns minValue).
//
@deftype UniformIntegerDist <IntegerDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
                setIntegerMin: (int) minValue setMax: (int) maxValue;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setIntegerMin: (int) minValue setMax: (int) maxValue;
USING
- setIntegerMin: (int) minValue setMax: (int) maxValue;
- (int) getIntegerMin;
- (int) getIntegerMax;
- (int) getIntegerWithMin: (int) minValue withMax: (int) maxValue;
@end
```

Documentation

```
Name:  | UniformInteger

Description:  | Distribution returning integers over an interval [min,max]

Algorithm:  | Transformation of uniform unsigneds [0,unsignedMax]

Reference:  | (none)

Parameters:  | integer minValue
| integer maxValue

State:  | unsigned long long int currentCount

Implementation  | Setting minValue=maxValue is allowed.
comment:        | Then minValue is retured "with probability 1".

UniformIntegerDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:
```

```
|  -getIntegerSample:                         7.534 uS
|  -average # of generator calls per variate: 1.000 (U)
|  -distribution net time:                    3.836 uS

Relative speed:
| Speed 0.491 (time 2.037) relative to MT19937 getUnsignedSample
--------------- . -----------------------------------------------------------
--
```

(./BernoulliDist.html)

(./Interface.html)

(./UniformUnsignedDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

UniformUnsignedDist Class

```
//
// UniformUnsigned --
//   A generator of non-negative integral values uniformly distributed across
//   a closed interval [min,max].  (The interval in-
cludes both its endpoints.)
//   Setting minValue == maxValue is allowed (and returns minValue).
//
@deftype UniformUnsignedDist <UnsignedDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
                setUnsignedMin: (unsigned) minValue setMax: (un-
signed) maxValue;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setUnsignedMin: (unsigned) minValue setMax: (un-
signed) maxValue;
USING
- setUnsignedMin: (unsigned) minValue setMax: (unsigned) maxValue;
- (unsigned) getUnsignedMin;
- (unsigned) getUnsignedMax;
- (unsigned) getUnsignedWithMin: (unsigned) minVal
           withMax: (unsigned) maxVal;
@end
```

Documentation

```
Name:  | UniformUnsigned

Description:  | Distribution returning unsigneds over an interval [min,max]

Algorithm:  | Transformation of uniform unsigned [0,unsignedMax]

Reference:  | (none)

Parameters:  | unsigned minValue
| unsigned maxValue

State:  | unsigned long long int currentCount

Implementation  | Setting minValue=maxValue is allowed.
```

```
comment:        | Then minValue is returned "with probability 1".

UniformUnsignedDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getUnsignedSample:                        7.632 uS
|  -average # of generator calls per variate: 1.000 (U)
|  -distribution net time:                    3.934 uS

Relative speed:
| Speed 0.485 (time 2.064) relative to MT19937 getUnsignedSample
--------------- . -------------------------------------------------------------
--
```

(./UniformIntegerDist.html)

(./Interface.html)

(./UniformDoubleDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

UniformDoubleDist Class

```
//
// UniformDouble --
//   A generator of floating point values uniformly distributed
//   across a half-open interval [min,max). (The interval includes
//   the lower endpoint but excludes the upper endpoint.)
//   NOTE: Setting minValue == maxValue is allowed (and returns minValue).
//
@deftype UniformDoubleDist <DoubleDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
                setDoubleMin: (double) minValue setMax: (double) maxValue;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setDoubleMin: (double) minValue setMax: (double) maxValue;
USING
- setDoubleMin: (double) minValue setMax: (double) maxValue;
- (double) getDoubleMin;
- (double) getDoubleMax;
- (double) getDoubleWithMin: (double) minValue withMax: (double) maxValue;
@end
```

Documentation

```
Name:  | UniformDouble

Description:  | Distribution returning doubles over an interval [min,max)

Algorithm:  | Transformation of uniform double [0.0,1.0)

Reference:  | (none)

Parameters:  | double minValue
| double maxValue

State:  | unsigned long long int currentCount

Implementation | Setting minValue=maxValue is allowed.
comment:  | Then minValue is returned "with probability 1".

UniformDoubleDist | On a 486/66, the following measures were obtained,
```

```
speed:  | using MT19937 as the underlying generator:

|  -getDoubleSample:                          13.777 uS
|  -average # of generator calls per variate:  1.000 (D)
|  -distribution net time:                     3.618 uS

Relative speed:
| Speed 0.268 (time 3.726) relative to MT19937 getUnsignedSample
--------------- . ------------------------------------------------------------
--
```

(./UniformUnsignedDist.html)

(./Interface.html)

(./ExponentialDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

ExponentialDistribution Class

```
//
// ExponentialDistribution -- a well-
known continuous probability distribution
//
@deftype ExponentialDist <DoubleDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
                setMean: (double) mean;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
                setMean: (double) mean;
USING
- setMean: (double) mean;
- (double) getMean;
- (double) getSampleWithMean: (double) mean ;
@end
```

Documentation

```
Name:  | ExponentialDistribution

Description:  | Exponential distribution returning double values

Algorithm:  | Transformation method

Reference:  | Russell 1992

Parameters:  | double Mean

State:  | unsigned long long int currentCount

ExponentialDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getDoubleSample:                          20.905 uS
|  -average # of generator calls per variate:  1.000 (D)
|  -distribution net time:                    10.746 uS

Relative speed:
| Speed 0.177 (time 5.653) relative to MT19937 getUnsignedSample
```

--------------- . ---
--

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-26

GammaDistribution Class

```
//
// GammaDistribution -- a well-known continuous probability distribution
//
@deftype GammaDist <DoubleDistribution, CREATABLE>
CREATING
+ create: aZone setGenerator: generator
setAlpha: (double) alpha setBeta: (double) beta;
+ create: aZone setGenerator: generator
setVirtualGenerator: (unsigned) vGen
setAlpha: (double) alpha setBeta: (double) beta;
USING
- setAlpha: (double) alpha setBeta: (double) beta;
- (double) getAlpha;
- (double) getBeta;
- (double) getSampleWithAlpha: (double) alpha
withBeta: (double) beta;
@end
```

Documentation

```
Name:  | GammaDistribution

Description:  | Gamma distribution returning double values

Algorithm:  | Transformation, rejection and xxx methods.

Reference:  | Watkins 1994.

Parameters:  | double Alpha
| double Beta

State:  | unsigned long long int currentCount

GammaDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getDoubleSample:                          119.424 uS
|  -average # of generator calls per variate:   6.224 (D)
|  -distribution net time:                     56.194 uS
```

```
Relative speed:
| 0.00310 (time 32.294) relative to MT19937 getUnsignedSample
--------------- . ------------------------------------------------------------
--
```

(./ExponentialDist.html)

(./Interface.html)

(./NormalDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Normal and LogNormalDistribution Classes

```
//
// NormalDistribution -- a well-known continuous probability distribution
// LogNormalDistribution -- also well-known ...
//
@deftype NormalDist    <Normal, CREATABLE> @end
@deftype LogNormalDist <Normal, CREATABLE> @end
```

NormalDist Documentation

```
Name:  | NormalDistribution

Description:  | Normal (Gaussian) distribution returning double values

Algorithm:  | The Box-Muller method

Reference:  | W.H. Press et al, "Numerical Recipes in C" (2nd ed.).
| Cambridge University Press, 1992.

Parameters:  | double Mean
| double Variance

State:  | double stored_value
| BOOL stored
| unsigned long long int currentCount

NormalDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getDoubleSample:                          24.729 uS
|  -average # of generator calls per variate:  1.273 (D)
|  -distribution net time:                    11.797 uS

Relative speed:
| Speed 0.150 (time 6.687) relative to MT19937 getUnsignedSample
--------------- . ----------------------------------------------------------
--
```

LogNormalDist Documentation

```
Name:  | LogNormalDistribution

Description:  | Log-Normal distribution returning double values

Reference:  | W.H. Press et al, "Numerical Recipes in C" (2nd. ed.).
| Cambridge University Press, 1992.

Algorithm:  | The Box-Muller method

Parameters:  | double Mean
| double Variance

State:  | double stored_value
| BOOL stored
| unsigned long long int currentCount

Output types:  | double

Output Quality:  | All bits should be safe.

LogNormalDist | On a 486/66, the following measures were obtained,
speed:  | using MT19937 as the underlying generator:

|  -getDoubleSample:                          33.430 uS
|  -average # of generator calls per variate:  1.273 (D)
|  -distribution net time:                    20.498 uS

Relative speed:
| Speed 0.111 (time 9.040) relative to MT19937 getUnsignedSample
--------------- . -----------------------------------------------------------
--
```

(GammaDist.html)

(Interface.html)

(RandomBitDist.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Usage Guide

Table Of Contents:

Overview

The random library basically consists of two parts, the actual pseudorandom number algorithms and the transformations that change the output of the pseudorandom number algorithms into the simulated distributions. The pseudorandom number algorithms are called generators throughout these documents and many times the "pseudo" is dropped from pseudorandom, even though there is no true random number generator implemented in Swarm.

Generator Usage Notes

Primary Author: Sven Thommesen
(mailto:sthomme@humsci.auburn.edu)
Swarm rel. 1.0.3 (random rel. 0.8)
Notes on the use of the random number generators
Sven Thommesen <sthomme@humsci.auburn.edu>
1997-09-01

```
IMPROVEMENTS over v. 0.6:

A host of new generators, located on the web or in the literature, have
been added since the last version of Random. There is now a total of 36
different generators defined! Some of these have immense periods, some
are very fast, and some have much better statistical properties than
the old generators.

A new *type* of generator, the 'split' generator, has been introduced
in the form of L'Ecuyer's C2LCGXgen and C4LCGXgen generators.

A 'split' generator is a long-period generator for which we are able to
split the period into arbitrary sub-periods, which we can access quickly.
```

We then configure the generator as having a number (A) of 'vir-
tual generators',
each of which can address a number (2^v) of sub-segments of length 2^w.
These parameters (A,v,w) are user selectable when the generator is created.
(As an example, for C4LCGXgen the default values are A=128, v=31, w=41.)
The advantage is that the subsegments act as statistically independent
streams of random numbers.

In addition to the -getUnsignedSample method, generators now also supply
floating point output in the range [0.0,1.0), in the form of these methods:

```
   -(float)       getFloatSample; // using 1 unsigned value
   -(double)      getThinDoubleSample; // using 1 unsigned value
   -(double)      getDoubleSample; // using 2 unsigned values
   -(long double) getLongDoubleSample; // using 2 unsigned values
```

The last method is not portable across architectures, since the length of a
long double varies between machines.

Generators may now be started with a single seed, *or* with a vector of seeds
whose length is generator dependent. (PMMLCG requires 1 integer for a
seed, while MT19937 needs 624 of them.)

Generators now remember what seed values they were started with. They also
count how many variates they have delivered (i.e., how many calls to
-getUnsignedSample they have serviced.)

There are a few arbitrary seed values, DEFAULTSEED, DEFAULT-
SEED1, DEFAULTSEED2,
DEFAULTSEED3, DEFAULTSEED4 defined. There is the value RANDOMSEED, which will
be different every time it is invoked. And there is value START-
SEED, which will
by default equal DEFAULTSEED, but will instead be equal to RANDOMSEED if you
start your program with the '-varySeed' command line parameter.

The generators have gained a new creation method, '+createWithDe-
faults: aZone',
which creates the generator and initializes it with STARTSEED. Split
generators get default values for A,v,w.

CHANGES since v. 0.6:

The generator classes have changed names to where they all end in '-gen'.
A simple search-and-replace in your code will get you up and running again.

(Or perhaps you'll want to try one of the new generators?)

A bug in SWBgen was corrected. Code for ACG and SCG was also changed.

The -verifySelf method is gone. Instead see the test program located in /random/testR0.

The 'getState:' method has been named 'putStateInto: (void *) buffer', and the 'setState:' method is now 'setStateFrom: (void *) buffer'. A quick search-and-replace fixes things in your code.

Note: these methods have also changed somewhat, as has the size of the data being saved. As a result, v. 0.7 generators will refuse to 'setStateFrom' data saved by v. 0.6 objects.

There should be fewer changes like this in the next release.

USAGE notes:

I shall first discuss 'simple' generators, and then discuss how 'split' generators differ from 'simple' ones.

In the following text, wherever I use PSWBgen as an example you may substitute any other generator.

Note: any name that starts with 'my' is meant to designate a variable which you have defined in your own program.

SIMPLE generators:

You create a generator in one of 3 ways:

(a) the lazy way:

```
   myGenerator = [ PSWBgen createWithDefaults: [self getZone] ];
```

This allocates the object and initializes it with STARTSEED.

(b) using a single seed value:

```
   myGenerator = [ PSWBgen create: [self getZone]
setStateFromSeed: mySeed ];
```

This allocates the object and initializes it with your seed value. If the object actually requires a vector of seed values to fill the state, this method generates the rest of the values needed using an inline PMMLCG generator.

You can find out later what seed value was used to initialize the generator:

```
   myUnsigned = [ myGenerator getInitialSeed ];
```

And you can find out what the largest valid seed value is by calling

```
   myUnsigned = [ myGenerator getMaxSeedValue ];
```

(For the generators defined in v. 0.7, this value is 2^32-1 for all of them. The seed may not be 0.)

You may reset the generator's state at any time using this method:

```
   [ myGenerator setStateFromSeed: mySeedValue ];
```

(c) using a vector of seed values:

Assume we have defined a fixed array at compile time:

```
   unsigned int mySeedVector [vectorLength];
```

Then we can do this:

```
   myGenerator = [ PSWBgen create: [self getZone]
setStateFromSeeds: mySeedVector ];
```

You can find out how many seed values are required by asking

```
   myUnsigned = [ myGenerator lengthOfSeedVector ];
```

(Obviously, you must first successfully have created the object to do this, for example using createWithDefaults!)

And we allocate the seed vector dynamically this way:

```
   unsigned int *mySeedVector;
   mySeedVector = [[self getZone] alloc: [ myGenerator lengthOfSeedVector]];
```

You can find out what vector of seed values was used to initial-
ize the object:

```
   unsigned int *myVector;
```

```
   myVector = [ myGenerator getInitialSeeds ];
```

And you can find out the largest seed values that are allowed for the particular generator:

```
   myVector = [ myGenerator getMaxSeedValues ];
```

(These values vary from generator to generator, and they may not be the same for all elements of the vector for a given generator. Seeds may not be 0.)

NOTE: in the above two calls, the variable myVector is set to point to an array internal to the generator. If you want to preserve the values, you need to allocate space in your program either statically or dynamically, and use a for-loop to copy data from myVector[i] to myAllocatedVector[i].

You may reset the generator's state at any time by using the method

```
   [ myGenerator setStateFromSeeds: (unsigned *) mySeedVector ];
```

This will also reset to 0 the currentCount variable.

NOTE: if you set a generator's state from a vector of seeds, the call

```
   myUnsignedValue = [ myGenerator getInitialSeed ];
```

will return a value of 0 (an invalid seed). On the other hand, if you initialize the generator with a single seed value, the call

```
   mySeedVector = [ myGenerator getInitialSeeds ];
```

will return the seed vector that would produce identical output to what you obtained using the single seed.

(d) antithetic values:

You can make the generator serve up antithetic values by setting:

```
   [ myGenerator setAntithetic: myBooleanValue ];
```

If set, this makes -getUnsignedSample return (unsignedMax-x) instead of x, and the floating point methods return (1.0 - y) instead of y.

You can ascertain if this flag is set by calling

```
   myBooleanValue = [ myGenerator getAntithetic ];
```

(e) generator output:

You obtain successive pseudorandom numbers from a generator by calling

```
   myUnsignedValue = [ myGenerator getUnsignedSample ];
```

The largest value thus returned can be found by asking

```
   myUnsignedValue = [ myGenerator getUnsignedMax ];
```

(The smallest value returned is always 0.)

If you would rather have floating point output in the range [0.0,1.0),
you call one of these:

```
// Using 1 unsigned value to fill the mantissa:

   myFloatValue  = [ myGenerator getFloatSample ];
   myDoubleValue = [ myGenerator getThinDoubleSample ];

// Using 2 unsigned values to fill the mantissa:

   myDoubleValue     = [ myGenerator getDoubleSample ];
   myLongDoubleValue = [ myGenerator getLongDoubleSample ];
```

NOTE that the last method is not portable, because the size of a long double
varies and hence the precision varies between architectures.

Finally, you can obtain a count of how many variates have been generated:

```
   myLongLongInt = [ myGenerator getCurrentCount ];
```

(currentCount is an unsigned long long int, which counts up to 2^64.)

SPLIT generators:

A 'split' generator requires us to specify the configuration (A,v,w)
at create time:

```
   myGenerator = [ C4LCGXgen create: [self getZone]
setA: 64 setv: 20 setw: 76
setStateFromSeed: mySeedValue ];
```

```
   myGenerator = [ C4LCGXgen create: [self getZone]
setA: 32 setv: 25 setw: 60
setStateFromSeeds: (unsigned *) mySeedVector ];

(In both cases, the only limitation is that A * 2^v * 2^w must be less than
the generator's period, 2^60 for C2LCGX and 2^120 for C4LCGX.)

For obtaining output, we need to specify which of the A 'virtual' generators
we want to draw from:

   myUnsignedValue = [ myGenerator getUnsignedSample: 12 ];
   myFloatValue    = [ myGenerator getFloatSample: myVirtualGenerator ];
   myDoubleValue   = [ myGenerator getThinDoubleSample: someUnsignedValue ];
   myDoubleValue   = [ myGenerator getDoubleSample: 32 ];
   myLongDoubleValue = [ myGenerator getLongDoubleSample: 0 ];

Note: virtual generators are numbered from 0 to A-1.

Obtaining the current count of variates generated likewise:

   myLongLongInt = [ myGenerator getCurrentCount: myVirtualGenerator ];
   myLongLongInt = [ myGenerator getCurrentSegment: myVirtualGenerator ];

The latter call indicates what segment number the virtual generator is
currently drawing numbers from.

Other than these methods, the other methods discussed above under 'simple'
generators are the same for 'split' generators.

In *addition* to this, 'split' generators have the following methods to
manage the virtual generators:

   // Place all virtual generators at the start of the first segment:
   [ myGenerator initAll ]; // done automatically at creation

   // Place all virtual generators back to the start of the current segment:
   [ myGenerator restartAll ];

   // Place all virtual generators at the start of the next segment:
   [ myGenerator advanceAll ];

   // Place all virtual generators at the start of the indicated segment:
   [ myGenerator jumpAllToSegment: myLongLongIntValue ];

You may also address individual virtual generators:

   [ myGenerator initGenerator:    myVgen ];
```

```
   [ myGenerator restartGenerator: myVgen ];
   [ myGenerator advanceGenerator: myVgen ];
   [ myGenerator jumpGenerator:    myVgen    toSegment: myLongLongIntValue ];
```

InternalState methods common to simple and split generators:

```
   // Print (most of) the object's state data to a stream:
   [ myNormalDist describe: myStream ];
```

The stream myStream may be created thus:

```
id myStream = [ OutStream create: [self getZone] setFileStream: stdout ]; or
id myStream = [ OutStream create: [self getZone] setFileStream: stderr ];

   // Get the (class) name of the object:
   myString = [ myNormalDist getName ];

   // Get the object's 'magic number', used by putStateInto / setStateFrom:
   myUnsigned = [ myNormalDist getMagic ];
```

SAVING AND RESETTING STATE:

You may save, and later restore, the internal state of a generator using these methods:

```
   // Get the size of the memory buffer needed by put-
StateInto / setStateFrom:
   myUnsigned = [ myGenerator getStateSize ];

   // Extract the generator's state data into your memory buffer:
   [ myGenerator putStateInto: myBuffer ];

   // Set the generator's state from data in a memory buffer:
   [ myGenerator setStateFrom: myBuffer ];
```

To illustrate, assume the following data definitions:

```
   FILE *myFile;
   const char *myFileName = "MyGenFile.bin"; // or whatever
   int stateSizeG;
   id stateBufG;
   int status;
```

The following code shows how to save an object’s state to disk:
(You should add your own code to deal with disk file errors,
either aborting or printing out error messages.)

```
   // Ask how big a buffer we need:
   stateSizeG = [ myGenerator getStateSize ];

   // Allocate memory for the buffer:
   stateBufG  = [[self getZone] alloc: stateSizeG];

   // Ask the generator to put state data into the buffer:
   [ myGenerator putStateInto: (void *) stateBufG ];

   // Open a disk file for output:
   myFile = fopen(myFileName, "w");
   if (myFile == NULL) { }; // error on open: disk full, or no permissions

   // Write the state buffer to disk in binary form:
   status = fwrite(stateBufG, stateSizeG, 1, myFile);
   if (status < 1) { }; // error on write: disk full?

   // Close the file
   status = fclose(myFile);
   if (status) { }; // error on close ?

   // Free the memory allocated to the buffer:
   [[self getZone] free: stateBufG];

   // Or, for test purposes, just zero the buffer data instead:
   // memset(stateBufG, 0, stateSizeG);
```

This code shows how to set an object’s state from a disk file:

```
   // Ask how big a buffer we need:
   stateSizeG = [ myGenerator getStateSize ];

   // Allocate memory for the buffer:
   stateBufG  = [[self getZone] alloc: stateSizeG];

   // Open a disk file for input:
   myFile = fopen(myFileName, "r");
   if (myFile == NULL) { }; // error on open: file not found

   // Read state data into the memory buffer:
   status = fread(stateBufG, stateSizeG, 1, myFile);
   if (status < 1) { }; // error on read
```

```
   // Close the file:
   status = fclose(myFile);
   if (status) { }; // error on close

   // Ask the generator set its state from the buffer data:
   [ myGenerator setStateFrom: (void *) stateBufG ];

   // Free the memory allocated to the buffer:
   [[self getZone] free: stateBufG];
```

TESTING GENERATORS:

Since v. 0.6 we have done some rudimentary statistical testing of the implemented generators, using Marsaglia's Diehard tests and the ENT tests.

The results of these tests are summarized in document /random/docs/doc.quality.generators, where test results as well as period length, state size and execution times are listed. You can use these data to select a generator that suits your simulation.

Some brief comments:

a) the tests show that old generators SCG and LCG are of poor quality and should be avoided.

b) the lagged-Fibonacci based generators (ACG, SWB, PSWB) all fail Diehard's 'Birthday spacings test', for reasons having to do with their lattice structure.
These generators are only conditionally recommended.

c) The rest of the 32-bit generators (i.e. generators that fill all 32 bits of an unsigned int) pass all tests, and are recommended at this time.

(Note that while a test may show that a generator is bad, passing a number of tests does not prove that a generator is good!)

d) The 31-bit generators all fail the same set of tests. Some of these cannot be passed by a generator whose output has a 'stuck' bit. Until I clear up with Prof. Marsaglia how to interpret these results, I believe all the 31-bit generators are in the 'recommended' category.

However, a cautionary note: while the PMMLCG generators pass the tests, they have a very short period (less than 2^31) and should only be used for 'toy' simulations. You don't want your generator(s) to 'go around' and start repeating themselves !

For what it's worth, Professor L'Ecuyer recommends his own C4LCGX and C2MRG3 generators as well as Matsumoto's TT800 (the monster MT19937 hadn't been released yet), and Prof. Marsaglia recommends his own Multiply-With-Carry generators (MWCA, MWCB, C3MWC, RWC2, RWC8="Mother").

UTILITY OBJECTS PROVIDED:

The following objects have been defined in <random/random.m>, and are immediately accessible from anywhere in your program:

```
id <MT19937>             randomGenerator;
id <UniformIntegerDist>   uniformIntRand;
id <UniformUnsignedDist>  uniformUnsRand;
id <UniformDoubleDist>    uniformDblRand;
```

The 3 distribution objects all draw their random numbers from the MT19937 generator, which has a period of 2^19937 (10^6001) and is quite fast.

Distribution Usage Notes

Primary Author: Sven Thommesen
(mailto:sthomme@humsci.auburn.edu)
Swarm rel. 1.0 (random rel. 0.6)
Notes on the use of the distribution objects
Sven Thommesen <sthomme@humsci.auburn.edu>
1997-09-01

IMPROVEMENTS over v. 0.6:

One new distribution class, BernoulliDist, has been added. It returns binary values (yes/true/1) with a given probability (while the old RandomBitDist has a fixed 50% probability, a fair coin toss.)

Distributions now have a new create method, '+createWithDefaults: aZone'. This method creates the distribution object, and also a new generator object for its exclusive use. Each distribution class has a different default generator class assigned. These generators are initialized with STARTSEED, which by default equals the fixed value DEFAULTSEED, but will be equal to the varying RANDOMSEED if you start your program with the command line parameter '-varySeed'.

All distributions have code to interact with the new 'split' generators.

UniformIntegerDist and UniformUnsignedDist now allow you to set parameter minValue equal to maxValue. In this case that value is returned every time.

UniformDoubleDist also allows this, even if the set [x,x) is mathematically suspect ...

NormalDist and LogNormalDist now allow you to specify zero Variance, in which case the values returned are the Mean and exp(Mean) respectively.

CHANGES since v. 0.6:

The distribution classes have changed names to where they all end in 'Dist'. A simple search-and-replace in your code will get you back up and running.

The strong distinction between 'frozen' and 'un-frozen' distribution objects in v. 0.6 has been softened considerably. You may now set and reset the default parameters as often as you wish, and you may make calls for variates with given parameters even if different default parameters have been set.

The generation of uniform(0,1) floating point values has been moved from the distribution objects into the generator objects. Thus, if all you need is a uniform(0,1) double, you have no need of a distribution but can get what you desire from a generator.

Note that the generators fill the mantissa of a double from two 32-bit unsigned values in a different manner from v. 0.6 distributions, so output will be a bit different in the new version.

A bug in LogNormalDist has been fixed.

The 'getState:' method has been named 'putStateInto: (void *) buffer', and the 'setState:' method is now 'setStateFrom: (void *) buffer. A quick search-and-replace fixes things in your code.

But note: these methods have also changed somewhat, as has the size of the data being saved. As a result, v. 0.7 distributions will refuse to 'setStateFrom' data saved by v. 0.6 objects.

There should be fewer changes like this in the next release.

USAGE notes:

Each distribution object must be assigned a random generator on creation. You may not re-assign generators to a distribution after it has been created. You may, however, connect more than one distribution object to one generator, so that they end up drawing output from the generator in an interleaved fashion. Or you may create a new generator for each distribution object.

Each distribution has its own set of (1-2) parameters. You may deal with these parameters in two different ways: (1) you assign a set of default parameter values to the object on creation, and draw from the distribution using those parameters. Or (2), you may refrain from assigning default parameters, in which case you must specify the (possibly different) desired parameters on each call. You can (re-)set the default parameters any time, and you may call for a variate with specified parameter values even if different default parameters have been set.

Just like the generator objects, the distribution objects allow you to save and later restore their internal state.

Where I use NormalDistribution in examples below, substitute any other distribution and its parameters as needed.

NOTE: any name that starts with 'my' is meant to designate a variable which you have defined in your own program.

You create a distribution in one of several ways:

(a) the lazy way:

```
   myNormalDist = [ NormalDist createWithDefaults: [self getZone]];
```

This method will create a distribution object with no default parameters set, as well as a fresh generator object connected to it. The generator object is initialized with STARTSEED (see the discussion above). Different distribution classes use different generators for this purpose.

(b) Without default parameters, using a simple generator:

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: mySimpleGenerator ];
```

'myGenerator' must of course first have been set to point to a random generator of the 'simple' type. Note that you cannot assign a different generator to a distribution after it has been created.

You can create the generator at the same time as the distribution:

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: [TT800gen create: [self getZone]
setStateFromSeed: 34453]         ];
```

(c) Without default parameters, using a split generator:

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: mySplitGenerator
setVirtualGenerator: 7 ];
```

or perhaps

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: [C4LCGXgen createWithDefaults: [self getZone]]
setVirtualGenerator: 99 ];
```

A split generator can be thought of as comprising a set of virtual generators (streams of random numbers), and a distribution object must be 'connected' to one of these streams. You cannot re-assign the generator or the virtual generator after a distribution object has been created.

In all these cases (a) - (c), when we want to obtain a random variate from this distribution object we need to specify the parameters:

```
   myDouble = [ myNormalDist getSampleWithMean: 3.3 withVariance: 1.7];
```

You can use different parameters for every call. (And you can use this call even if default parameters have been set.)

(d) With default parameters, using a simple generator:

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: mySimpleGenerator
setMean: 7.6 setVariance: 1.2 ];
```

(e) With default parameters, using a split generator:

```
   myNormalDist = [ NormalDist create: [self getZone]
setGenerator: mySplitGenerator
setVirtualGenerator: 33
setMean: 3.2 setVariance: 2.1 ];
```

In these cases, we do not need to specify parameters to get a random number:

```
   myDouble = [ myNormalDist getDoubleSample ];
```

However, you *are* allowed to specify parameters even if default parameters have been set.

(Of course, different distributions have different parameters: RandomBitDist has none, the Uniform objects have minimum and maximum limit val-
ues, NormalDist
and LogNormalDist use Mean and Variance, ExponentialDist only Mean, and GammaDist used alpha and beta. See the file random/distributions.h for the specific methods available.)

(f) You may reset the default parameters this way, as often as you like:

```
   [ myNormalDist setMean: 3.3 setVariance: 2.2 ];
```

(g) You can obtain the current values of parameters:

```
   // Default parameters:
   myDouble1 = [ myNormalDist getMean ];
   myDouble2 = [ myNormalDist getVariance ];
   myDouble3 = [ myNormalDist getStdDev ];

   // Get a pointer to the generator object:
   myOtherGenerator = [ myNormalDist getGenerator ];

   // Get the number of the virtual generator (if a split generator is used):
   myUnsignedValue  = [ myNormalDist getVirtualGenerator];

   // Find out if default parameters have been set:
   myBoolean        = [ myNormalDist getOptionsInitialized ];

   // Find out how many variates the object has delivered so far:
   // (The counter is an unsigned long long int, which goes up to 2^64.)
   myLongLongInt    = [ myNormalDist getCurrentCount ];
```

(h) You can reset the variate counter and other state variables this way:

```
   [ myNormalDist reset ];
```

   This is most likely done in conjunction with resetting the connected generator, using [myGenerator setStateFromSeed: mySeedValue];

(i) Finally, we have the InternalState protocol methods:

```
   // Print (most of) the object's state data to a stream:
   [ myNormalDist describe: myStream ];
```

The stream myStream may be created thus:
id myStream = [OutStream create: [self getZone] setFileStream: stdout]; or
id myStream = [OutStream create: [self getZone] setFileStream: stderr];

```
   // Get the (class) name of the object:
   myString = [ myNormalDist getName ];

   // Get the object's 'magic number', used by putStateInto / setStateFrom:
   myUnsigned = [ myNormalDist getMagic ];
```

You may save, and later restore, the internal state of a distribution object using these methods:

```
   // Get the size of the memory buffer needed by put-
StateInto / setStateFrom:
   myUnsigned = [ myNormalDist getStateSize ];

   // Extract the distribution's state data into your memory buffer:
   [ myNormalDist putStateInto: myBuffer ];

   // Set the distribution's state from data in a memory buffer:
   [ myNormalDist setStateFrom: myBuffer ];
```

To illustrate, assume the following data definitions:

```
   FILE *myFile;
   const char *myFileName = "MyDistFile.bin"; // or whatever
   int stateSizeD;
   id stateBufD;
   int status;
```

The following code shows how to save an object's state to disk:
(You should add your own code to deal with disk file errors,
either aborting or printing out error messages.)

```
   // Ask how big a buffer we need:
   stateSizeD = [ myNormalDist getStateSize ];

   // Allocate memory for the buffer:
   stateBufD  = [[self getZone] alloc: stateSizeD];

   // Ask the distribution object to put state data into the buffer:
```

```
[ myNormalDist putStateInto: (void *) stateBufD ];

// Open a disk file for output:
myFile = fopen(myFileName, "w");
if (myFile == NULL) { }; // error on open: disk full, or no permissions

// Write the state buffer to disk in binary form:
status = fwrite(stateBufD, stateSizeD, 1, myFile);
if (status < 1) { }; // error on write: disk full?

// Close the file
status = fclose(myFile);
if (status) { }; // error on close ?

// Free the memory allocated to the buffer:
[[self getZone] free: stateBufD];

// Or, for test purposes, just zero the buffer data instead:
// memset(stateBufD, 0, stateSizeD);
```

This code shows how to set an object's state from a disk file:

```
// Ask how big a buffer we need:
stateSizeD = [ myNormalDist getStateSize ];

// Allocate memory for the buffer:
stateBufD  = [[self getZone] alloc: stateSizeD];

// Open a disk file for input:
myFile = fopen(myFileName, "r");
if (myFile == NULL) { }; // error on open: file not found

// Read state data into the memory buffer:
status = fread(stateBufD, stateSizeD, 1, myFile);
if (status < 1) { }; // error on read

// Close the file:
status = fclose(myFile);
if (status) { }; // error on close

// Ask the distribution object to set its state from the buffer data:
[ myNormalDist setStateFrom: (void *) stateBufD ];

// Free the memory allocated to the buffer:
[[self getZone] free: stateBufD];
```

UTILITY OBJECTS PROVIDED:

The following objects have been defined in <random/random.m>, and are immediately accessible from anywhere in your program:

```
id <MT19937>             randomGenerator;
id <UniformIntegerDist>   uniformIntRand;
id <UniformUnsignedDist>  uniformUnsRand;
id <UniformDoubleDist>    uniformDblRand;
```

The 3 distribution objects all draw their random numbers from the MT19937 generator, which has a period of 2^19937 (10^6001) and is quite fast.

(Interface.html)

(../random.html)

(AdvancedUsage.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Advanced Usage Guide

There should also be a general expository on the use of pseudo random number generation in simulation, especially as applied to Swarm.

(./UsageGuide.html)

(../random.html)

(./InterfaceDesign.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Interface Design Notes

Never has a library interface been debated as this one was (and still is). This document will attempt to provide some of the rational for the decisions made by echoing some of the controversial issues covered in the debates on the interface. (These discussions overlapped into implementation. But, the point, here, is to give people an idea of the complexity of the issues involved in useful pseudo-random number use in simulation.)

Issues:

- Different classes for generators or variants? Ref: http://www.sela.co.il:8080/swtrain/new/shai/tgp.ps
- Include non-recommended generators?
- Seed vs. State for generators?
- What is "state" and how do we handle it?
- Larger issue of debugging support via an object reporting its own configuration via a special method and how this relates to object state reporting and object saving.
- Separate classes for separate distribution types (e.g. double, integer, unsigned :: continuous/discrete)?
- Problems with the transformation from generator output to distribution output.
- What precision should be supported?
- Should the generator underlying a distribution be visible through the distribution?
- Efficiency!
- What distributions to provide?
- What generators to provide?
- Open vs. Closed intervals?
- Programmed default generators, seeds, and states?
- Programming by committee? (a.k.a. Software process!)
- Shorthand create methods?
- Potential phase optimizations?
- Random number streams, syncing limit cycles, and proper simulation vs. code debugging. (This is a HOT ONE.)
- Library interchangeability.

These are just the questions and issues. Some addition should follow (when there's time) addressing each issue and giving rationale for what we chose. But, that's a big job.

(AdvancedUsage.html)

(../random.html)

(Implementation.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Random Library: Implementation Notes

The generators supplied with this release have been subjected to statistical testing using George Marsaglia's Diehard tests as well as John Walker's entropy tests (ENT). The results of these tests are summarized in the document `$(SWARMHOME)/src/random/docs/doc.quality.generators`. Other properties of the generators are summarized in `$(SWARMHOME)/src/random/docs/generators.table`. Some notes on how to choose a generator for your simulation are found in `$(SWARMHOME)/src/random/docs/CHOOSING.A.GENERATOR`.

The ENT test is included in the tarball of test programs found at the Swarm web site. The Diehard tests are copyright and hence are not, but they can be downloaded from the web at http://www.hku.hk/internet/randomCD.html (http://www.hku.hk/internet/randomCD.html)
.

The distribution objects have not been statistically tested.

(./InterfaceDesign.html)

(../random.html)

(../random.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Simtools Library

Documentation and Implementation Status

Compatibility

- 1.0.4 -> 1.0.5: GUISwarm now inherits from GUIComposite. Because GUIComposite handles the passthru of archiving keys to tkobjc primitives, the method setControlPanelGeometryRecordName is no longer needed; instead, the macro SET_WINDOW_GEOMETRY_RECORD_NAME is provided.

- 1.0.3 -> 1.0.4: All functions maintain backward compatibility. There are additional features, however, and features previously undocumented.

Overview

Simtools is the stdlib of Swarm... In other words, it is the library where we have parked many miscellaneous classes which, while very useful, are not specific enough to be placed in any other library. So, for example, simtools contains InFile, ObjectLoader and other I/O classes, the ControlPanel and GUISwarm which support GUI interactions, but also UName which simply generates unique names for objects.

Documentation

The *objectbase* library is defined and documented following the basic structure of the Library Interface Conventions

(defobj/library.html)
. The following standard sections of documentation are available:

- Interface Reference

(simtools/Interface.html)

- Usage Guide
(simtools/UsageGuide.html)
- Advanced Usage Notes
- Interface Design Notes
- Implementation Notes

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Simtools Library: Documentation and Implementation Status

The simtools library has undergone an upgrade to the status of a library as of 1.0.4. This means the interface now conforms to the library interfaces specifications, the format of the documentation now reflects this. Other than new features, there should be no affect on the user.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Simtools Library: Interface Reference

The *simtools* library uses the module definition conventions
(../defobj/module.html)
established by the *defobj* library
(../defobj.html)
. It also depends on standard supertypes and classes defined by this library.

Dependencies

Following are the other header files imported by <simtools.h>:

```
#import <objectbase.h>
#import <random.h> // for global RNGs...
#import <tkobjc.h> // some methods return graphics stuff...
```

The *objectbase* library interface is included to provide the basic object support. The *random* library is included to provide support for default generators and distributions, the use of which is deprecated, but are included for backwards compatibility.

Special global functions - global.h

Users need to include simtools.h in their code in order to call `initSwarm()` but also in order to get access to a set of important pre-initialized objects which are generated in every simulation (e.g. `uniformIntRand`, `probeDisplayManager`, and so on).

void initSwarm(int argc, char ** argv)
: Initializes the Swarm libraries. This should be the call in any Swarm code you write. `argc` and `argv` are the input parameters to `main()`.

ProbeDisplayManager * probeDisplayManager
: A global instance of the ProbeDisplayManager object. You can use this object to create new probe displays, and should schedule an update message to it in your display schedule so as to update the probe displays it has created. See the documentation on probes (../probing.html) for more details.

int swarmGUIMode
: A flag that indicates whether graphics are enabled. Set automatically based on whether the switch `-batchmode` was present as one of the command line arguments starting the application (e.g. `heatbugs -batchmode`). If it is false, then

globalTkInterp (from libtkobjc) is set to be a normal Tcl interpreter instead of a Tcl/Tk/BLT interpreter.

Object Types (with links to reference pages)

As simtools is a collection of tools that are only loosely related to each other - the class hierarchy is virtually flat. Most of the classes are standalone.

Contents:

(NSelect.html)

- UName
(UName.html)
- ActiveGraph
(ActiveGraph.html)
- ActiveOutFile
(ActiveOutFile.html)

(../simtools.html)

(../simtools.html)

(UsageGuide.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

Name

GUISwarm

Synopsis

```
//
// GUISwarm -- a version of the Swarm class which is graphics aware
//

@protocol GUISwarm <SwarmProcess>
@end
```

Description

GUISwarm is a subclass of Swarm that is used as a toplevel Swarm for simulations running with a graphical user interface. The GUISwarm creates a ControlPanel automatically for you and defines a "go" method that interprets the state of the ControlPanel to keep things running in response to user input. Users subclass GUISwarm much like they subclass a normal Swarm, implementing the same kind of buildObjects, buildActions, and activateIn methods. When you are done building your Observer Swarm, start it as a toplevel via `[myGUISwarm go]`.

The control panel places a few responsibilities on the GUISwarm subclass author. In particular, a message to `[controlPanel doTkEvents]` should be scheduled fairly frequently - only when that method is executed will the user interface update (and the stop button be checked). Also, it is often useful to use `[controlPanel setStateStopped]` to wait for the user to indicate they're ready for execution to proceed.

- go

  Start the activity running, and also handle user requests via the control panel. Returns either `Completed` (the model ran until requested to terminate) or `ControlStateQuit` (the user pressed the quit button).

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ControlPanel - class to control the top level SwarmProcess

Synopsis

```
// ControlPanel --
// an class used to control the top-level SwarmProcess. It
// accepts manipulations both by the Swarm it's controlling and
// the ActionCache.
//

@protocol ControlPanel <SwarmObject>
CREATING
- createEnd;
USING
- getState;
- setState: s;
- waitForControlEvent;
- startInActivity: activityID;
- setStateRunning;
- setStateStopped;
- setStateStepping;
- setStateQuit;
- setStateNextTime;

// Deprecated methods
- doTkEvents;
- getPanel;
@end

// State Symbols for the ControlPanel.
extern id ControlStateRunning, ControlStateStopped;
extern id ControlStateStepping, ControlStateNextTime, ControlStateQuit;
```

Description

ControlPanel keeps track of the users requests to run, stop, quit, or time step the simulation. It cooperates with the GUISwarm to control the execution of activities in Swarm.

- getState

Get the current button state of the controlpanel. Is one of ControlStateRunning, ControlStateStopped, ControlStateStepping, ControlStateNextTime, or ControlStateQuit.

- setStateRunning, - setStateStopped, - setStateStepping, - setStateQuit, - setStateNextTime

  Sets the state to the indicated value, and also sends the appropriate stop or terminate message to the running activity tree. The - setStateStopped message is particularly useful since it will cause the simulation to stop until the user interactively sets it back in motion (in other words, this method is useful in generating a software-triggered pause).

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ActionCache - a class to manage threads and Swarms

Synopsis

```
//
// ActionCache -
//
// A class that provides a smart bag into which actions can be
// thrown by other threads and Swarms intended for insertion on
// it's Swarm's schedule.
//

@protocol ActionCache <SwarmObject>
CREATING
- setControlPanel: cp;
- setScheduleContext: context;
- setWindowGeometryRecordName: (const char *)name;
- createEnd;

USING
- insertAction: actionHolder;
- deliverActions;
- sendActionOfType: (id <Symbol>)type toExecute: (const char *)cmd;
- sendStartAction;
- sendStopAction;
- sendStepAction;
- sendNextAction;
- sendQuitAction;
- verifyActions;

- createProcCtrl;
- getPanel;
- doTkEvents; //should change to pollGUI or something
@end

// Type Symbols for ActionCache
extern id <Symbol> Control, Probing, Spatial;
// Error Symbols for ActionCache
extern id <Symbol> InvalidActionType, ActionTypeNotImplemented;
```

Description

- setWindowGeometryRecordName: (const char *)name
  Set the name to be used for archiving the geometry of the control panel widget.
  There is no default; the position of the control panel will not be preserved
  unless you use this method.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ProbeDisplay - a class to display ProbeMaps

Synopsis

```
//
// ProbeDisplay
//
// A class which generates a GUI to a ProbeMap of probes applied to a
// given target object.
//
@protocol ProbeDisplay <SwarmObject>
CREATING
- setProbedObject: anObject;
- setProbeMap: (ProbeMap *)probeMap;
- setWindowGeometryRecordName: (const char *)name;
USING
- getProbedObject;
- getProbeMap;
- update; // implemented in SimpleProbeDisplay...
@end
```

Description

Create Phase Protocol

- setProbedObject: anObject

  This method *must* be called.

- setProbeMap: (ProbeMap *)probeMap

  This is an optional create phase method - if no probeMap is specified the ProbeDisplay will ask the probedObject for a ProbeMap using the getProbeMap method described below... The default behaviour of this method will be to return the probeLibrary's copy of the probeMap for the class of the target object.

- setWindowGeometryRecordName: (const char *)name

In order to record the window position of a ProbeDisplay, it is necessary to give it a name. This method is used to give an instance ProbeDisplay a name, which will used by the Archiver when recording its geometry information.

Regular Protocol

- getProbedObject
- getProbeMap
- update

  This method maintains consistency between the values of the probedObject's variables and the values which are displayed in the ProbeDisplay. Ideally, this method should be called every time the object is modified by the simulation. In practice, the user schedules an update on the probeDisplayManager which in turn communicates to all the active ProbeDisplays in the system.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

CompleteProbeDisplay

Synopsis

```
//
// CompleteProbeDisplay --
// A class which generates a GUI to a complete ProbeMap of probes applied
// to a given target object (by complete we mean that all the probes for
// the target object's class and its superclasses are included).
//

@protocol CompleteProbeDisplay <SwarmObject>
CREATING
- setProbedObject: anObject;
- setWindowGeometryRecordName: (const char *)name;
USING
- getProbedObject;
- update;
- getMarkedForDropFlag;
@end
```

Description

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ProbeDisplayManager

Synopsis

```
//
// ProbeDisplayManager --
// A singleton class whose instance is used to manage all the
// ProbeDisplays created by the user during a GUI run of the
// simulation.
//

@protocol ProbeDisplayManager <SwarmObject>
USING
- createProbeDisplayFor: anObject;
- createProbeDisplayFor: anObject setWindowGeometryRecord-
Name: (const char *)name;

- createCompleteProbeDisplayFor: anObject;
- createCompleteProbeDisplayFor: anObject setWindowGeometryRecord-
Name: (const char *)name;

- addProbeDisplay: probeDisplay;
- removeProbeDisplayFor: anObject;
- removeProbeDisplay: probeDisplay;
- setDropImmediatelyFlag: (BOOL)dropImmediateFlag;
- update;
@end
```

Description

- addProbeDisplay: pd
- removeProbeDisplay: pd
  These methods are used to specify which Probe Displays should be managed by the ProbeDisplayManager.
- update
  This method will recursively send an update message to all the Probe Displays managed by the ProbeDisplayManager.

- createProbeDisplayFor: anObject
- createDefaultProbeDisplayFor: anObject
- createCompleteProbeDisplayFor: anObject
- createProbeDisplayFor: anObject setWindowGeometryRecordName: (const char *)name
- createCompleteProbeDisplayFor: anObject setWindowGeometryRecordName: (const char *)name;
- dropProbeDisplaysFor: anObject

  These methods allow a one-step creation/removal of ProbeDisplays onto a given object (dropProbeDisplaysFor: will actually *drop* the appropriate ProbeDisplays). The setWindowGeometryRecordName: creation variety can be used to set names for archiving geometry information.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1997-12-19

Name

InFile - class to perform file input

Synopsis

```
//
// InFile --
// A class which was intended to support file input. There have been
// justified requests from our userbase to re-design this interface.
//

@protocol InFile <SwarmObject>
CREATING
+ create: aZone withName: (const char *)theName;
USING
- (int)getWord: (char *)aWord;
- (int)getLine: (char *)aLine;
- (int)getInt: (int *)anInt;
- (int)getUnsigned: (unsigned *)anUnsigned;
- (int)getLong: (long *)aLong;
- (int)getUnsignedLong: (unsigned long *)anUnsLong;
- (int)getDouble: (double *)aDouble;
- (int)getFloat: (float *)aFloat;
- (int)getChar: (char *)aChar;
- (int)unGetChar: (char)aChar;
- (int)skipLine;
@end
```

Description

This class is intended to simplify the input file-I/O in Swarm. It essentially deals with the detailed file opening and closing routines thus alleviating the need for C file I/O procedure calls.

+ create: aZone withName: (const char *)theName
  This is the create method for InFiles, where theName is, of course the name of the file to open.

- (int)getWord: (char *)aWord

We use the term word instead of string because the returned value will not contain spaces, tabs and newlines -> thus getWord...

- (int)getLine: (char *)aLine
  This routine loads the argument string with the characters up to, but not including a newline character.
- (int)getInt: (int *)anInt;
- (int)getUnsigned: (unsigned *)anUnsigned;
- (int)getLong: (long *)aLong;
- (int)getUnsignedLong: (unsigned long *)anUnsLong;
- (int)getDouble: (double *)aDouble;
- (int)getFloat: (float *)aFloat;
- (int)getChar: (char *)aChar;
  These routines take a pointer to the required type, and read into it an instance of that type from the open file. In case of failure these message returns 0.
- (int)unGetChar: (char) aChar
- (int)skipLine
  These routines are used by ObjectLoader, but are documented in case the user wishes to define his/her own file format which may require backtracking or skipping lines.
- (void)drop
  This method must be called when the user wants to close the file. Only by dropping the InFile object will the file truly be closed!

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

OutFile - class to perform file output

Synopsis

```
//
// OutFile --
// A class which was intended to support file output.
//

@protocol OutFile <SwarmObject>
CREATING
+ create: aZone withName: (const char *)theName;
USING
- putString: (const char *)aString;
- putInt: (int) anInt ;
- putUnsigned: (unsigned) anUnsigned;
- putLong: (long) aLong ;
- putUnsignedLong: (unsigned long) anUnsLong ;
- putDouble: (double) aDouble ;
- putFloat: (float) aFloat ;
- putChar: (char) aChar ;
- putTab ;
- putNewLine ;
@end
```

Description

This class is intended to simplify output file-I/O in Swarm. It essentially deals with the detailed file opening and closing routines thus alleviating the need for C file I/O procedure calls.

+ create: aZone withName: (const char *)theName
  This is the create method for OutFiles, where theName is, of course the name of the file to open.

- putString: (const char *)aString

- putInt: (int) anInt

- putUnsigned: (unsigned) anUnsigned;
- putLong: (long) aLong;
- putUnsignedLong: (unsigned long)anUnsLong;
- putDouble: (double)aDouble
- putFloat: (float)aFloat
- putChar: (char)aChar
- putTab: (char)aChar
- putNewLine
  These routines take an instance of the appropriate type and write it into the open file. The return value is simply self since no error detection is performed.
- (void)drop
  This method must be called when the user wants to close the file. Only by dropping the OutFile object will the file truly be closed!

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

AppendFile - class for appended file output

Synopsis

```
//
// AppendFile --
// A class to support appended file output.
//
@protocol AppendFile <SwarmObject>CREATING
+ create: aZone withName: (const char *)theName;
USING
- putString: (const char *)aString;
- putInt: (int)anInt;
- putUnsigned: (unsigned)anUnsigned;
- putLong: (long)aLong;
- putUnsignedLong: (unsigned long)anUnsLong;
- putDouble: (double)aDouble;
- putFloat: (float)aFloat;
- putChar: (char)aChar;
- putTab;
- putNewLine;
@end
```

Description

This class subclasses from OutFile, the only difference in functionality, being that it opens a given file in *Append*

mode, rather than *Overwrite*mode, which OutFile does.

+ create: aZone withName: (const char *)theName
  This is the create method for AppendFiles, where theName is, of course the name of the file to open.

- putString: (const char *)aString

- putInt: (int)anInt

- putUnsigned: (unsigned)anUnsigned;

- putLong: (long)aLong;
- putUnsignedLong: (unsigned long)anUnsLong;
- putDouble: (double)aDouble
- putFloat: (float)aFloat
- putChar: (char)aChar
- putTab: (char)aChar
- putNewLine

  These routines take an instance of the appropriate type and write it into the open file. The return value is simply self since no error detection is performed.
- (void)drop

  This method *must* be called when the user wants to close the file. Only by dropping the AppendFile object will the file truly be closed!

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1997-12-19

Name

ObjectLoader - class to load objects IVars from a file

Synopsis

```
//
// ObjectLoader --
// An attempt to design some form of standard for object loading.
//

@protocol ObjectLoader <SwarmObject>
USING
+ load: anObject from: aFileObject;
+ load: anObject fromFileNamed: (const char *)aFileName;

- setFileObject: aFileObject;
- loadObject: anObject;
// In case the same class is being loaded multiple times...
- updateCache: exampleTarget;
@end
```

Description

This class is used to initialize the variables of a target object from a data file. The data file is required to have a

very simple format:

```
# A Comment

@begin

aVarName aValue
aVarName aValue # A comment
aVarName aValue
#
# A Comment
#
aVarName aValue
```

@end

Note: A single file can be used to initialize multiple objects by simply inserting more than one @begin @end
block.

+ load: anObject fromFileNamed: (const char *)aFileName This convenience class method will load anObject from the file named aFileName (the class will open the file, initialize the object with its contents and then close the file).
+ load: anObject from: aFileObject This convenience class method will load anObject from the previously opened aFileObject without returning an actual instance of the ObjectLoader class (also, the FileObject remains open after the method has been called).
- setFileObject: aFileObject If the user bothers to generates an actual instance of the ObjectLoader class, s/he can set the source fileObject which the instance should use by sending it this message.
- loadObject: anObject This is the message the user must send an instance of the ObjectLoader class in order to initialize the target object from the requested file.
- updateCache This method should be called if an ObjectLoader instance is going to initialize a large number of objects from the same class.
- setTemplateProbeMap: aProbeMap This method is used to specify which variables of the target object(s) should be loaded by the ObjectLoader instance to which this message was sent.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ObjectSaver - class to save object's IVars to a file

Synopsis

```
//
// ObjectSaver --
// An attempt to design some form of standard for object saving.
// The ProbeMap argument is used to specify what variables get saved.
//

@protocol ObjectSaver <SwarmObject>
USING
+ save: anObject to: aFileObject;
+ save: anObject to: aFileObject withTemplate: aProbeMap;
+ save: anObject toFileNamed: (const char *)aFileName;
+ save: anObject toFileNamed: (const char *)aFileName withTemplate: (Probe-
Map *)aProbeMap;

- setFileObject: aFileObject;
- setTemplateProbeMap: aProbeMap;
- saveObject: anObject;
@end
```

Description

This class is used to write an object's variables to a specified file. If only a subset of the variables should be

written out, the set is specified by a template ProbeMap (where the ProbeMap will contain Probes for those variables which should be saved).

+ save: anObject to: aFileObject

+ save: anObject to: aFileObject withTemplate: aProbeMap

+ save: anObject toFileNamed: (const char *)aFileName

+ save: anObject toFileNamed: (const char *)aFileName withTemplate: aProbeMap

As in the case of ObjectLoader we provide a set of convience class methods which will save the entire target object (or the subset of variables specified in a template) without actually returning an instance of ObjectSaver to the user.

- setFileObject: aFileObject
  If the user bothers to generates an actual instance of the ObjectSaver class, s/he can set the target fileObject which the instance should use by sending it this message.
- saveObject: anObject
  This is the message the user must send an instance of the ObjectSaver class in order to save the state of the target object into the requested file.
- setTemplateProbeMap: aProbeMap
  This method is used to specify which variables of the source object(s) should be saved by the ObjectSaver instance to which this message was sent.

Saving floats and doubles: The desired number of significant figures for all floats and doubles can be set in the global instance of the ProbeLibrary. However, if notset, it will default to a value of six (6) significant figures, which is identical to previous versions that did not have this feature. See the documentation on probes (../probing.html)
for how to set this precision.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

QSort - class to sort a collection

Synopsis

```
//
// QSort --
// A class (not to be instantiated) wrapper around the C sort routine.
//

@protocol QSort <SwarmObject>
USING
+ sortObjectsIn: aCollection;
+ (void)sortObjectsIn: aCollection using: (SEL)aSelector;
+ (void)sortNumbersIn: aCollection;
+ (void)sortNumbersIn: aCollection us-
ing: (int(*)(const void*,const void*)) comp_fun;
+ (void)reverseOrderOf: aCollection;
@end
```

Description

```
QSort is simply a "wrapper" for the C native "qsort" function, as applied to a Swarm collection.
The values will appear in ascending order by default.  Reversing the order of a collection
can be made by calling reverseOrderOf.  All these methods modify the underlying collection,
so any indexes should always be regenerated.

    +(void) sortObjectsIn:  aCollection
          Given a collection, this method will sort the objects with the "compare" function
          for the object.  If the objects don't provide a compare function, it uses
          the default from the defobj library.

    +(void)sortObjectsIn:  aCollection using:  (SEL)aSelector
          This method allows the user to specify a given message selector to use to
          "compare" the objects.

    +(void) sortNumbersIn:  aCollection;
          This method is intended specifically for integer comparisons, again using
          a default "compare" function.  The default assumes numbers should be monotonically
          increasing.
```

+(void) sortNumbersIn: aCollection using: (int(*)(const void*,const void*))comp_fun
As above, except that a specific function is given.

+(void) reverseOrderOf: aCollection
Reverses the current order of a collection. To make a "reversed" (i.e. decending) sort, simply call one of the appropriate "sort..." methods on a collection, then call this function on the same collection.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

NSelect - class to select exactly N elements at random from a collection

---

Synopsis

```
//
// NSelect --
// A class (not to be instantiated) wrapper around a Knuth algorithm
// for the selection of exactly N elements form a collection without
// repetition.
//

@protocol NSelect <SwarmObject>
USING
+ (void)select: (int)n from: aCollection into: bCollection;
@end
```

Description

NSelect selects exactly N elements from a collection without repetition. A target collection must be provided.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

UName - class to generate unique names for objects/agents

Synopsis

```
//
// UName --
// A class used to generate unique names (e.g. "critter1", "critter2" etc).
//

@protocol UName <SwarmObject>
CREATING
+ create: aZone setBaseName: (const char *)aString;
+ create: aZone setBaseNameObject: aStringObject;

- setBaseName: (const char *)aString;
- setBaseNameObject: aStringObject;
USING
- (const char *)getNewName;
- getNewNameObject;
@end
```

Description

This class is used to generate unique names (agent0, agent1, agent2...) for objects in a simulation. The user will typically create an instance of the UName class initialized with a baseName presented either as a (const char*) or an object of class String

(../collections/String.html)
. The user can then request new names, again either as (const char *)'s or as instances of the String class. The user can also reset the counter used to generate the names in case s/he wants to restart naming objects with the same baseName.

> Note: Both in the case of initialization by (const char *) and initialization by an instance of the String class, the original is copied *not stored internally* so it is up to the user to free the original (const char *) or String instance if/when necessary!

+ create: aZone setBaseName: (const char *)aString

```
+ create: aZone setBaseNameObject: aStringObject

- setBaseName: (const char *)aString

- setBaseNameObject: aStringObject
      These messages should be used to set or reset the base name.
      Note: these methods will automatically reset the counter.

- (const char *)getNewName

- getNewNameObject
      Generates a new name in the requested format.

- resetCounter
      Resets the counter used as a suffix in the unique names generated.
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1997-12-19

Name

ActiveGraph - provides continous data feed between Swarm and the GUI

Synopsis

```
//
// ActiveGraph --
// This class provides the continuous data feed between Swarm and the GUI.
//

@protocol ActiveGraph <MessageProbe>
USING
- setElement: ge;
- setDataFeed: d;
- step;
@end
```

Description

```
An active graph object is the glue between a MessageProbe (for reading data) and a BLTGraph
GraphElement.  ActiveGraphs are created and told where to get data from and send it to, and
then are scheduled to actually do graphic functions.  This class is used by EZGraph, and we
expect to see less direct usage of it by end-users as more analysis tools (such as EZGraph)
internalize its functionality.

  - setElement:  get
       Sets the graph element used to draw on

  - setDataFeed:  d
       Sets the object that will be probed for data.

  - setProbedSelector:  (SEL)s
       Sets the message sent to the object for graphing (inherited from MessageProbe).

  - step
       Fires the probe, reads the value from the object, and draws it on the graph
       element.  The X value is implicitly the current simulation time.  Y is the
       value read.
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Name

ActiveGraph - provides continous data feed between Swarm and a file

Synopsis

```
//
// ActiveOutFile --
// Provides the continuous data feed between Swarm and a File.
//

@protocol ActiveOutFile <MessageProbe>
USING
- setFileObject: aFileObj;
- setDataFeed: d;
- step;
@end
```

Description

```
This is the file I/O equivalent of ActiveGraph:  it takes an OutFile object, a target (datafeed)
object, and a selector, which it uses to extract data from the object and send it to the file.
This class is used by EZGraph, and we expect to see less direct usage of it by end-users as
more analysis tools (such as EZGraph) internalize its functionality.

     - setFileObject:  anOutFile
           Sets the file object to which the data will be sent.

     - setDataFeed:  d
           Sets the object that will be probed for data.

     - setProbedSelector:  (SEL) s
           Sets the message sent to the target object (inherited from MessageProbe).

     - step
           Fires the probe, reads the value from the object, and sends it to the file.
           The X value is implicitly the current simulation time.  Y is the value read.
```

Marcus G. Daniels <mgd@santafe.edu>

(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Simtools Library: Usage Guide

Overview

The *simtools* library contains a number of semi-related tools that are useful for writing simulations.

(Interface.html)

(../simtools.html)

(../simtools.html)

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-10

TkObjc Documentation

Tcl/Tk

(ftp://ftp.smli.com/pub/tcl/)
is a scripting language and graphical widget set. TkObjc is a library of wrapper classes around Tk and BLT widgets. It's purpose is to provide a simple graphical interface while hiding most Tk-specific code from library users. To create and use graphical widgets, the user merely needs to create and use objects. Many of the objects here are straightforward wrappings of Tk widgets, but some (ButtonPanel, for instance) are combinations of other widgets, and others (Raster) are novel code.

TkObjc is currently based on Tcl 7.5, Tk 4.1
(ftp://ftp.smli.com/pub/tcl/)
, BLT 2.1
(ftp://ftp.neosoft.com/pub/tcl/NEW/)
or BLT 2.3
(http://www.tcltk.com/blt/index.html)
,and tclobjc-1.2
(ftp://ftp.santafe.edu/pub/swarm/needed-software/tclobjc-1.2.tar.gz)
(available from Swarm authors). Very little of this code is library-version dependent, however, as most of it works by directly calling the Tk interpreter.

The basic purpose of tkobjc is to package Tk functionality. Therefore, tkobjc's behaviour is similar to the Tk toolkit. For simple usage one should be able to get fairly far just by looking at this document, the header files, and the Swarm examples: more complicated graphical output will require the programmer have some familiarity with Tk.

---

globalTkInterp

TkObjc gets most of its work done by talking to the Tcl interpreter object `globalTkInterp`, an instance of class TkExtra. Strings handed to the interpreter's eval: method are evaluated by the Tcl interpreter and any side effects are performed. The libtclobjc library provides the functionality that glues Tk to Objective C. Users with simple needs never need to talk to the Tk interpreter directly.

Widget Base Class

All graphical widgets inherit from the Widget base class. Widget defines most of the behaviour needed: Widgets are created by the user with a particular parent, and then "pack"ed in order to draw them on the screen. All widgets have three essential things: a widget name used when running Tcl code, an Objective C name when sending messages from Tcl to those objects, and a parent.

+ createParent: (Widget *)p and - setParent: (Widget *)p
: When a widget is created it needs to be given a parent. The parent widget will be the widget's containing window. If no parent is given (ie, a parent of nil), then a toplevel Frame will be allocated automatically

- setWindowTitle: (const char *)s
: Sets the window manager title of the top level window

- pack, - packWith: (const char *)c
: Roughly, packing a widget makes it draw on the screen. The Tk packer allows complicated options to control widget layout: you can pass options via -packWith:. See documentation on Tk to learn more about packing details.

- getWidth, - setWidth: (int)n, - getHeight, - setHeight: (int)n, - getPositionX, - getPositionY, - setPositionX: (int)n Y: (int)n, etc.
: These methods allow you to control the position and size of the widget.

- (const char *)getWidgetName, -(const char *)getObjcName
: Get various names for the widgets. The WidgetName is a Tcl name for the window: the ObjcName is the name of the object itself. These are used if you choose to send strings off to Tcl directly for evaluation.

Simple Widgets

There are a variety of simple widgets which do straightforward things.

Frame

Frames are boxes other widgets fit in. They correspond to the Tk "frame" and "toplevel" widgets. Frames can be new windows, or subwindows in an existing window. You only need to create frames yourself if building complicated composite widgets: by default, a frame will be built automatically for widgets without parents.

Label

Labels are text strings used for displaying short information.

- setText: (const char *)t
: set the label's text string and resize it to fit.

Button

Buttons a user can click. Each button has a label and a command to be executed when the button is pressed.

- setText: (const char *)t
  set the button's text string and resize to fit.
- setCommand: (const char *)c
  set the button's command, what to do. This is a Tcl string passed off to the interpreter. See the ControlPanel code in simtools for an example.

Input Widgets

All widgets which take input from the user inherit from a base class, InputWidget. InputWidgets get their input in one of two ways: by being readable, or by being linked to a C variable.

InputWidget (base class)

- (const char *)getValue
  read the value as a string
- setValue: (const char *)v
  set the value as a string
- linkVariable: (void *)p Type: (int)type
  Link the widget directly to a C variable - any changes from the user will be automatically stored in the variable. Types are TCL_LINK_INT, TCL_LINK_BOOLEAN, TCL_LINK_DOUBLE, and TCL_LINK_STRING.

Entry

Handles data of any input type - user types in a string.

CheckButton

Boolean data only, click to set.

Compound Widgets

Some TkObjc widgets consist of several other widgets bound together into one object.

ButtonPanel

Several buttons bound together in one frame.

```
- addButtonName: (const char *)n Command: (const char *)c
     add a new button to the panel with the given name and command.
```

Data Display Widgets

TkObjc has two basic classes for displaying numerical data.

BLTGraph

A time series graph tool, based on BLT's graph widget. BLTGraph currently implements just a basic graph with multiple datasets, but should eventually support scaling and scrolling. For each BLTGraph you create one or many GraphElements, one per dataset to plot. GraphElements can be configured for appearance, and data can be added to the element to draw.

```
- title: (const char *)t, - axisLabelsX: (const char *)xl Y: (const char *)yl
     set labels and title on the graph.
- setScaleModeX: (int)xd Y: (int)ys
     takes two boolean values - whether to autoscale every timestep or instead
     to jump scale.
- setRangesXMin: (double)xmin Max: (double)xmax YMin: (double)ymin Max: (double)ymax
     sets the ranges for the graph, turns off autoscaling entirely
- (GraphElement *)createElement
     builds a new GraphElement for you to plot data with
- destroyElement: (GraphElement *)g
     remove the given graph element, destroy the plot.
```

GraphElements are where the real work are, where you give it data to plot. One BLTGraph can have multiple GraphElements.

```
- setLabel: (const char *)label
     set the label for this element
```

- setColor: (const char *)color, - setWidth: (int)w, - setDashes: (int)i, - setSymbol: (const char *)s
  change the apperance of the graph. The argument to dashes specifies the dash width. Valid symbols are line (default), square, circle, diamond, plus, cross, splus, scross.
- addX: (double)x Y: (double)y
  add the datapoint (x,y) to the graph
- resetData
  remove all data from this element, start over
- -(BLTVector *)getXData, -(BLTVector *)getYData
  Get the data in the graph as a BLTVector (undocumented)

Histo

Histogram display tool, based on BLT's barchart. The number of bins is fixed at creation time, then the user hands the Histo an array of datapoints (double or int) to display (or optionally an array of datapoints and *locations* where the bars should be drawn (specified as doubles).

- title: (const char *)t, - axisLabelsX: (const char *)xl Y: (const char *)yl
  title the histogram, give it axis labels.
- setNumPoints: (int)n Labels: (const char * const*) l Colors: (const char * const *)c
  initialize the histogram, tell it how big its dataset is. Labels and Colors are arrays of strings (one string per point) for text labels and the colours of the bars (the last two arguments are optional).
- drawHistoWithDouble: (double *)points, - drawHistoWithInt: (int *)points
  an array of points (two possible types) to draw in the histogram.
- drawHistoWithDouble: (double *)points atLocation: (double *)locations, - drawHistoWithInt: (int *)points atLocation: (double *)locations
  An array of points (two possible types) to draw in the histogram, and an array of specific locations in which to draw them (the bar width at every location will be of width 1.0).

Graphical Output

TkObjc has one basic class (and a subclass) for displaying graphical data. The model used here is a 2d raster display: an array of display elements. Display elements are typically simply pixels, blocks of a solid colour (specified via an XColormap). Other options are a pixmap (XPixmap) or arbitrary Xlib display code (implementation of XDrawer).

The display widgets used here do not have Tk counterparts: they are raw Xlib code written for flexibility and (relative) speed.

XColormap

Mechanism used to map numbers in the range [0, 255] to colour names. Create an XColormap, allocate colours in it, and pass it to a Raster widget for drawing.

- (BOOL)setColor: (Color)c ToName: (const char *)colorName
  set the specified color to the given colorname.
-(BOOL) setColor: (Color) c ToGrey: (double) g
  set the specified color to the grey value in [0, 1.0].
-(BOOL) setColor: (Color) c ToRed: (double) r Green: (double) g Blue: (double) b
  set the specified color to the red, green, blue vector in the range [0, 1.0].

Raster

2 dimensional, colour pixel images. Raster is based on a Tk frame widget with our own code for fast display of images. You can draw coloured dots on a Raster, or generic <XDrawer>s (not documented). Raster widgets are double buffered - the pixels you draw are not actually put on the screen until - drawSelf is called. In addition, Rasters handle mouse clicks.

- erase
  erase the window to black
- drawSelf
  render the Raster on the screen (complete the double buffer)
- setColorMap: (XColorMap *)colormap
  sets the colourmap used to translate logical colors to screen colors.
-(XColormap *) getColormap;
  returns the id of the colormap
- drawPointX: (int)x Y: (int)y Color: (Color)c;
  draw a point at the given coordinates with the given color.
- fillRectangleX0: (int)x0 Y0: (int)y0 X1: (int)x1 Y1: (int)y1 Color: (Color)c
  fill a whole rectangle with color
- setButton: (int)n Client: c Message: (SEL)s
  when button n (ButtonLeft, ButtonMiddle, or ButtonRight) is pressed, send the given message to the specified client object.
- (Display *)getDisplay
  returns the id of the display
- handleButton: (int)n X: (long)x Y: (long)y
  fields the button clicks

ZoomRaster

ZoomRaster is a subclass of Raster that implements a zoomable image. It handles translation between logical coordinates and screen coordinates.

- increaseZoom, - decreaseZoom, - setZoomFactor: (int)z
  change the zoom size
- (unsigned)getZoomFactor
  returns an integer representing the zoom size
- handleConfigureWidth: (unsigned)newWidth Height: (unsigned)newHeight
  reconfigures zoom when the window is resized

XDrawer, XPixmap

Using XDrawer and XPixmap, you can write arbitrary Xlib code to draw in the window. XDrawer is just a protocol for the drawOn:X:Y:GC:Caller: message.

XPixmap provides packaging for drawing pixmaps onto displays.

- setFile: (const char *)filename
  designates the file from which to read the pixmap image
- createEnd
  loads the pixmap and sets the display and window
- -(Pixmap) getPixmap
  returns the id of the pixmap data structure
- -(BOOL) getMasked
  returns YES if a mask has been set, NO if it hasn't
- -(Pixmap) getMask
  returns the id of the pixmap mask
- -(unsigned) getWidth, -(unsigned) getHeight
  returns the width or height of the pixmap

Canvas

There are the beginnings of a packaging of the Tk canvas widget that allows arbitrary drawing of complicated shapes. This code is still under development and is not thoroughly documented. Canvas

is a subclass of Widget and so inherits those methods but overrides createEnd.

CanvasItem

This is the superclass of any item that might be drawn on a Canvas, e.g. the geometric shapes, textItems, etc. This is an abstract class.

- setCanvas: canvas
  designates the id of the Canvas in which this item resides
- setTargetId: target
  designates the object to which this item refers
- setClickSel: (SEL)sel
  sets the message that will be effected upon a click on this item
- setMoveSel: (SEL)sel
  sets the message that will effect the motion of the item on the canvas
- setPostMoveSel: (SEL)sel
  sets the message that will dictate what happens after the item is moved
- createItem
  passthrough method to be implemented by subclasses
- createBindings
  establishes the bindings for the buttons: button 3 registers a "click", button 1 up and down sets the focus, and button 1 held down initiates motion.
- createEnd
  calls createItem and createBindings
- clicked
  sends the clickselector message to the referent
- initiateMoveX: (long)delta_x Y: (long)delta_y
  prepares for movement of the item within the canvas

Circle : CanvasItem

- setX: (int)x Y: (int)y
  set the x,y coordinates for the center of the circle
- setRadius: (int)r
  set the radius of the circle

Line : CanvasItem

- setTX: (int)tx TY: (int)ty LX: (int)lx LY: (int)ly
  set the endpoints of the line

Rectangle : CanvasItem

- setTX: (int)tx TY: (int)ty LX: (int)lx LY: (int)ly set the two diagonal corner coordinates of the rectangle

TextItem : CanvasItem

- setX: (int)x Y: (int)y
  sets the coordinate for the center of the text
- setText: (const char *)text
  sets the content of the text

CompositeItem

A CompositItem is simply a CanvasItem that consists of several pieces. (Why doesn't this inherit from CanvasItem?) The interface is identical to CanvasItem except for the additional method moveX:Y. Again, this is an abstract class.

- moveX: (long)delta_x Y: (long)delta_y

  passthrough method to be implemented by subclass

NodeItem : CompositeItem

- setX: (int)x Y: (int)y
  position the item on the canvas
- (int)getX, - (int)getY;
  return the x or y coordinate for the item
- setString: (const char *)string
  designate the string that will label this node
- setColor: (const char *)aColor
  designate the fill color of the node
- setBorderColor: (const char *)aColor
  designate the border color of the node
- setBorderWidth: (int)aVal
  set the width of the border in pixels

LinkItem : CompositeItem

- setFrom: from
  designate the object (node) that will be the source of the link
- setTo: to
  designate the node that will be the destination of the link
- setColor: (const char *)aColor
  set the color of the link

```
  - update
        redraw the link (esp due to motion of nodes)
```

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Analysis

This is the library where tools primarily related to analysis tasks, reside. This includes tools which simplify the task of graphing values or displaying distributions as well as more specific measurement tools (such as Average, Entropy).

Averager

Averager objects read a value (via a MessageProbe) from a collection (typically a list) of objects and collect statistics over them.

- setCollection: aCollection
  sets the collection of objects that will be probed.
- setProbedSelector: (SEL)s
  The message to be sent to each element in the collection (inherited from MessageProbe)
- (double)getAverage, - (double)getTotal, - (double)getMax, - (double)getMin - (int)getCount
  various statistics that the averager collects. These are read out of the object, not computed everytime they are asked for.
- update
  Update the average - send the message throughout the collection and update the average, total, etc. Schedule this message when you need to collect data.

Entropy

Entropy objects read probabilities (via a MessageProbe) from a collection of objects and calculate the entropy of the underlying distribution.

NOTE: This implies that the objects which are being polled must somehow calculate a valid probability, in other words, they are probably aware of a global property such as N (the number of agents in the system).

- setCollection: aCollection
  sets the collection of objects that will be probed.
- setProbedSelector: (SEL)s
  The message to be sent to each element in the collection (inherited from MessageProbe)
- (double)getEntropy
  Returns the calculated Entropy (i.e. the entropy is read out of the object, not computed everytime it is requested).

- update

    Update the entropy - send the message throughout the collection and update the entropy. Schedule this message when you need to collect the data.

EZGraph

This class allows the user to easily create graphs of various quantities in the model s/he is investigating. The user first creates the EZGraph, and then creates "Sequences"; (lines) which will appear in the graph. The sequences are generated based on data provided by a a single object or a collection of target objects, in reponse to a specified selector. So, for example, the following code generates a graph of statistics about the unhappiness of a list of heatbugs:

```
// Create the graph widget to display unhappiness.
  unhappyGraph = [EZGraph createBegin: [self getZone]];
  [unhappyGraph setTitle: "Unhappiness of bugs vs. time"];
  [unhappyGraph setAxisLabelsX: "time" Y: "unhappiness"];
  unhappyGraph = [unhappyGraph createEnd] ;

  [unhappyGraph createAverageSequence: "average unhappiness"
                         withFeedFrom: [heatbugModelSwarm getHeatbugList]
                          andSelector: M(getUnhappiness)] ;

  [unhappyGraph createMinSequence: "minimum unhappiness"
                     withFeedFrom: [heatbugModelSwarm getHeatbugList]
                      andSelector: M(getUnhappiness)] ;

  [unhappyGraph createMaxSequence: "maximum unhappiness"
                     withFeedFrom: [heatbugModelSwarm getHeatbugList]
                      andSelector: M(getUnhappiness)] ;
```

One of the features of the EZGraph is that it will automatically generate average, total, min, max and count sequences without the user having to mess with Averagers amd other low-level classes.

- setGraphics: (int)state

    If the user does not require a graphical display of the graph, s/he should call this method with state set to 0. The default state is 1 meaning that by default the data appears graphically in a window.

- setFileOutput: (int)state

    If the user would like the data for the sequences to be sent to a file, s/he should call this method with state set to 1. The default state is 0 meaning that by default, no file I/O is carried out by the EZGraph class.

- setTitle: (const char *)aTitle

The title is used in the graphical version of the EZGraph and appears at the top of the graph window.

- setAxisLabelsX: (const char *)xl Y: (const char *)yl

  Sets the labels on the graph. Only relevent for the graphical version.

- createSequence: (const char *)aName withFeedFrom: anObj andSelector: (SEL)aSel

  This creates a sequence in the EZGraph based on the return value provided by the object *anObj* when sent the selector *aSel*. If file I/O is enabled then the data will be sent to a file with the name *aName*, otherwise the *aName* argument is simply used as the legend for the graph element generated by EZGraph.

- createAverageSequence: (const char *)aName withFeedFrom: aCollection andSelector: (SEL)aSel
- createTotalSequence: (const char *)aName withFeedFrom: aCollection andSelector: (SEL)aSel
- createMinSequence: (const char *)aName withFeedFrom: aCollection andSelector: (SEL)aSel
- createMaxSequence: (const char *)aName withFeedFrom: aCollection andSelector: (SEL)aSel
- createCountSequence: (const char *)aName withFeedFrom: aCollection andSelector: (SEL)aSel

  These methods take a collection of objects and generate a sequence based on a statistic over the responses of the *entire* object-set (as opposed to a single object, as in - createSequence:withFeedFrom:andSelector:).

- step

  This method will cause the graphical display and/or the output files to be updated, i.e. the selector is applied to the target object(s) and the results are then displayed graphically and/or written to the appropriate output files.

- getGraph

  This method lets the user access the BltGraph generated internally by the EZGraph (assuming the creation of a graphical display was in fact required).

EZBin

This class allows the user to easily histogram data generated by a collection of objects. In addition the class will generate some standard statistics over the resulting dataset.

- setGraphics: (int)state

  If the user does not require a graphical display of the histogram, s/he should call this method with state set to 0. The default state is 1 meaning that by default the data appears graphically in a window.

- setFileOutput: (int)state
  If the user would like the data for the histogram to be sent to a file, s/he should call this method with state set to 1. The default state is 0 meaning that by default, no file I/O is carried out by the EZBin class.
- setTitle: (const char *)aTitle
  The title is used both in the graphical version of EZBin as well as in the file I/O where it serves as the name of the output file (when graphics are enabled it appears at the top of the histogram window).
- setAxisLabelsX: (const char *)xl Y: (const char *)yl
  Sets the labels on the histogram. Only relevent for the graphical version.
- setBinNum: (int)theBinNum
  Sets the number of bins the histogram will have.
- setLowerBound: (double)theMin
  Sets the lower bound on the histogram range.
- setUpperBound: (double)theMax
  Sets the upper bound on the histogram range.
- (int)getBinNum
- (double)getUpperBound
- (double)getLowerBound
  These methods retrieve the appropriate settings.
- setCollection: aCollection
  Set the collection of target objects which will be requested to generate the dataset for the histogram.
- setProbedSelector: (SEL)aSel
  Set the selector which will be applied to the objects in the specified collection in order to generate the dataset (inherited from *MessageProbe*).
- reset
  This method resets the histogram.
- update
  This method increases the amount of data included in the final dataset. In other words, it is possible to poll the same collection of objects repetitively before generating output (either graphically or on file).
- output
  This methods causes the graphical display to be updated (such that it takes into account the information extracted by the latest calls to update). When file I/O is enabled the number of entries per bin are sent to the output file.
- (int *)getDistribution
  Returns an array of integers (containing the number of entries which landed in each bin of the histogram).
- (int)getCount
  Get the number of entries which landed within the bounds of the histogram.

- (int)getOutliers
  Get the number of entries which landed out of the bounds of the histogram. Pressing the "o" key on the graphical representation of the histogram will display this value (both as an integer and as a percentage of the total number of attempted entries).
- (double)getMin
- (double)getMax
- (double)getAverage
- (double)getStd
  These messages are used to get some standard statistics over the dataset.

EZDistribution

This is a subclass of EZBin which normalizes the data and treats it as a distribution. This means that in addition to the statistics it can calculate by virtue of being a subclass of EZBin, it can also calculate the entropy of the distribution as well as return the probabilities associated with the individual bins.

- (double *)getProbabilities
  Returns an array of doubles representing the probability of every bin in the distribution.
- (double)getEntropy
  Returns the entropy of the distribution.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08

Probing and Displaying the Contents of a SwarmObject

Contents:

Overview

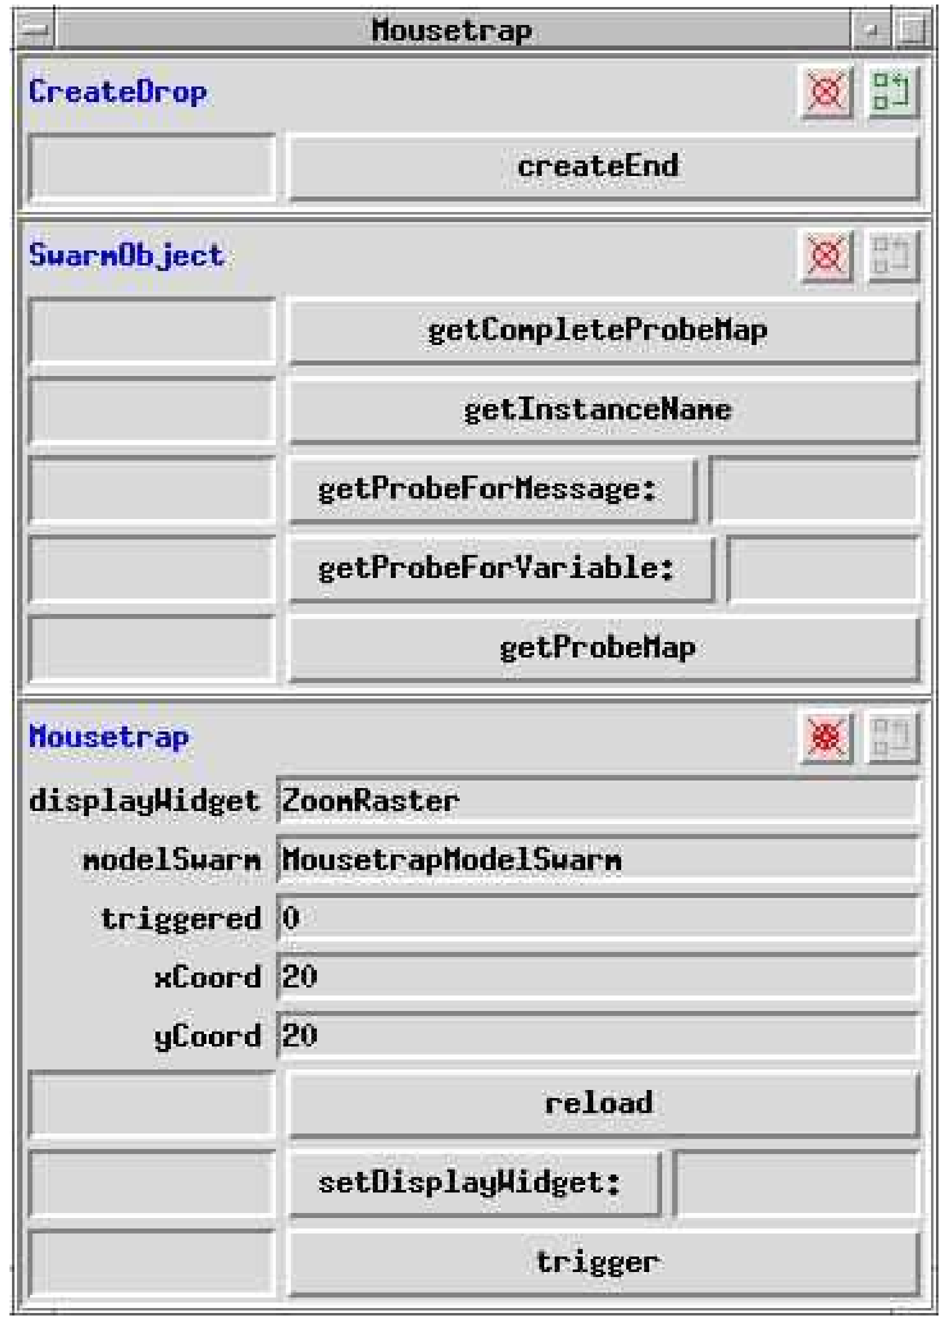
Mousetrap
CreateDrop
createEnd
SwarmObject
getCompleteProbeMap
getInstanceName
getProbeForMessage:
getProbeForVariable:
getProbeMap
Mousetrap
displayWidget ZoomRaster
modelSwarm MousetrapModelSwarm
triggered 0
xCoord 20
yCoord 20
reload
setDisplayWidget:
trigger

The Framework

Probes, ProbeMaps and ProbeDisplays allow the user to dynamically interact with the objects in their simulation. By *interaction* we mean reading/setting instance variables as well as dynamically generating method calls. The main point being that these interactions are not hardwired into the program code, but occur due to user-generated requests, mainly through the provided GUI.

The key to this capability is the Probe: in general a probe takes an object and either extracts the value of a specific variable, or calls a specific method. For this purpose we provide two subclasses: VarProbe and MessageProbe.

There are two main uses for probes: they can be fed into data-collection objects and serve as interfaces to the objects about which data is being collected (thus keeping the data-collection objects as general as possible) - the Averager class, for example, directly subclasses MessageProbe. Or, they can be used in order to generate a GUI to the individual objects in the simulation (the more common usage).

In order to generate a graphical version of a given probe, the programmer must place it inside a ProbeDisplay, which will automagically generate a window with the appropriate interface. Since, more often than not, the programmer will want to generate windows with more than one variable, we have designed the ProbeDisplay to deal with ProbeMaps rather than individual probes. Thus, by generating a ProbeMap containing exactly the right probes, a user is able to customise the window generated by the ProbeDisplay.

In order to facilitate the creation of all these different objects, the Swarm kernel provides some alternative methods for Probe/Map/Display generation:

- Direct Generation - which leaves the user in charge of all the details.
- Object Generation (Probes/ProbeMaps) - where we use method calls (defined in SwarmObject) on a given target object to generate the Probe/ProbeMap for that object. This allows the object to be aware that it is being probed and to control directly the contents of the Probe/ProbeMap being requested...
- Library Generation (Probes/ProbeMaps) - where the programmer "checks out" a unique, shared copy of a Probe/ProbeMap from the probeLibrary object (of class ProbeLibrary) provided by the kernel. By shared we mean that a similar request made at a different point in the code, will return a reference to the very same probe instance. *Note that, by default, Object Generation is equivalent to Library Generation since the default behaviour of SwarmObjects, when requested to create a ProbeMap, is to return a reference to the probeLibrary's unique copy!!!*
- ProbeDisplayManager Generation (ProbeDisplays) - where the programmer generates a ProbeDisplay directly, by requesting it from the probeDisplayManager object (of class ProbeDisplay) provided by the kernel (in graphics mode). The probeDisplayManager will create the ProbeDisplay based on a ProbeMap given to it by the probed object. Note: since the probeDisplayManager is only created in graphics mode, it is documented in the simtools area of the documentation. Here we emphasise simply that in order to generate a ProbeDisplay onto an object the programmer need only write the following line of code:

```
[probeDisplayManager createProbeDisplayFor: anObject] ;
```

And that the contents of the returned ProbeDisplay will be based on a ProbeMap obtained from the probed object itself.

Probe support resident in the objectbase *(objectbase.html)* library is documented with that library. Extra description, above and beyond the bare interface description for precision probing and SwarmObject is below:

Precision setting for floating-point numbers with probes

Global setting of precision

There are two types of global precision setting via ProbeLibrary (objectbase/ProbeLibrary.html) :

-setDisplayPrecision: (int) nSigFigsSaved
: Sets the number of significant figures for floating point (and double-floating) numbers diplayed on a GUI widget. Currently this is only implemented for VarProbes. The display uses the %*g sprintf-style formatting, which can vary slightly from implementation to implementation. If you set the number of significant figures to 3, then a float of value of 0.6344346 is displayed as 0.634 on the GUI widget. Note that this in no way affects the underlying stored value of the floating point number.

-setSavedPrecision: (int) nSigFigsSaved
: Sets the global default for the saving of floats through ObjectSaver. All objects with floats and doubles as instance variables are saved with the precision specified by this method. This is independent of the displayed precision of the same instance variable on a GUI widget.

To actually initialise these defaults: in the top level swarm, you should add the calls to the global ProbeLibrary instance (which is actually created by the initSwarm call in main) during the - createBegin method (this sets the precision in the global instance, before any probes are checked out of the instance. If neither method is called on probeLibrary, then the precision defaults to six significant figures in both cases.

Example: Here's how it's done in HeatbugObserverSwarm.m:

```
+ createBegin: aZone
{
  HeatbugObserverSwarm * obj;
```

```
    id <ProbeMap> probeMap;

    [...]

    probeMap = [EmptyProbeMap createBegin: aZone];
    [probeMap setProbedClass: [self class]];
    probeMap = [probeMap createEnd];
    // set the display defaults
    [probeLibrary setDisplayPrecision: 3];
    // typically saved precision would be higher than displayed precision
    // for statistical and data analysis purposes
    [probeLibrary setSavedPrecision: 10];

    // Add in a bunch of variables, one per simulation parameters
    [probeMap addProbe: [probeLibrary getProbeForVariable: "displayFrequency"
      inClass: [self class]]];
    [...]

    // Now install our custom probeMap into the probeLibrary.
    [probeLibrary setProbeMap: probeMap For: [self class]];
    return obj;
  }
```

Setting precision for individual probes

The formatting for an individual probe can be set directly, using a sprintf-style formatting string. Typically, customProbeMaps are created in the + createBegin factory method for a Swarm or a SwarmObject. To set the formatting for a floating point probe, the method from VarProbe

(objectbase/VarProbe.html)
is used:

```
-setFormatFloat:  (const char *)floatFormat
      is applied to the instance of the VarProbe "checked-out" of the global probe
      library instance.  The sprintf-formatting string can "over-ride" the "%g"
      format set by the global precision (as above) (Typically "%g" chooses between
      the "%f" and "%e", depending on the size of the exponent - which is implementation-dependent
      - this method allows you to explicitly set the type of display).
```

In the following example, it is desired that the number of significant figures for the floating point variable randomMoveProbability is three (3). Currently (Swarm 1.0.5) this is only works for VarProbes and not MessageProbes, as yet.

Example Here's how it's done in HeatbugModelSwarm.m:

```
+ createBegin: aZone
{
```

```
  HeatbugModelSwarm * obj;
  id <ProbeMap> probeMap;
  id floatProbe;

  [...]
  // the -setFloatFormat is applied to the probe which is
  // "returned" from the call to probeLibrary
  floatProbe = [[probeLibrary getProbeForVariable: "randomMoveProbability"
    inClass: [self class]]
                                  setFloatFormat: ".3f"];

  // now we have the probe - put it back into the customMap
  [probeMap addProbe: floatProbe];

  [...]
  return obj;
}
```

Or, more compactly:

```
+ createBegin: aZone
{
  HeatbugModelSwarm *obj;
  id <ProbeMap> probeMap;

  [...]

  [probeMap addProbe: [[probeLibrary getProbeForVari-
able: "randomMoveProbability"
    inClass: [self class]]
                             setFloatFormat: "%.3f"]];
  [...]
  return obj;
}
```

SwarmObject Support for Probing

As mentioned earlier, all SwarmObjects can generate Probes/ProbeMaps. By default, they achieve this by "checking out" the appropriate references from the probeLibrary.

Interface definition
(objectbase/SwarmObject.html)
from objectbase
(objectbase.html)
library.

```
  - (const char*)getInstanceName

        If you want to put a specific title on the ProbeDisplay window associated
        with your object, simply override this method.  By default it returns the
        class name of the object.  This is why the default ProbeDisplay window title
        is always the class name of the probed object.
```

```
Note:  SwarmObjects do not actually have an instance variable called "instanceName"
- so if you want to store a name for your object, you will need to provide your own
storage at the subclass level and rewrite getInstanceName to refer to this variable.
```

```
  - getProbeForVariable:  (const char *)aVariable

  - getProbeForMessage:  (const char *)aMessage

  - getProbeMap

  - getCompleteProbeMap
```

Probe support resident in the *simtools* library:

ProbeDisplay

```
The ProbeDisplay class is meant to serve a dual role, represented by the two example windows
on this page:
```

On the one hand if an object to be probed is specified *without* any particular ProbeMap being specified, then the ProbeDisplay generated will provide a window of class DefaultProbeDisplay, which displays all the variables resident in that class structure.

On the other hand, if a ProbeMap *is* specified then the ProbeDisplay follows exactly the specification as represented by the contents of the ProbeMap (above). When used in this manner, ProbeDisplays can generate tailored interfaces to objects (so for example, we have purposefully hidden certain instance variables in the HeatbugModelSwarm class, and have shown only one of the methods which the class understands).

Graphical Interface Usage

Common to both the standard ProbeDisplay and the CompleteProbeDisplay:

- The different fields in the ProbeDisplay can be updated by typing in new values and pressing Return. However, certain fields (containing pointers or ids, for example) cannot be modified and will generate a beep if such a modification is attempted.
- If an instance variable/argument slot is defined to hold an object, then that object can be drag&dropped into another variable/argument slot by clicking on it with the first mouse button (a small rectangle with the name of the object will appear - simply drag it to another object-typed variable/argument

slot and release the mouse button).

- Also, if an instance variable/argument slot is defined to hold an object, then that object can be inspected by clicking the entry for that variable/argument slot with the third mouse button(a ProbeDisplay for that object will be generated).

Available only on the standard ProbeDisplay:

- Note that the sunken label at the top of the ProbeDisplay is also active. By clicking on it with the first mouse button you get a drag&drop'able representation of *self*. By clicking on it with the third mouse button you get a CompleteProbeDisplay to *self*.

Available only on the ProbeDisplay:

- The green "superclass" button can be used to display the succesive superclasses of the object being probed.
- The red "hide" button can be used to hide classes which are irrelevant thus reducing clutter.
- The hide button on the lowest class in the hierarchy has a special meaning since clicking on it dismisses the entire ProbeDisplay.

Interface definition

(simtools/ProbeDisplay.html)
from simtools
(simtools.html)
library.

ProbeDisplayManager

The normal Swarm simulation will probably only ever contain one instance of this class, namely the probeDisplayManager. This object is used for automatic generation of ProbeDisplays: given an object it will attempt to get a ProbeMap from the object. If the object does not respond to the method - getProbeMap (i.e. it does not inherit from SwarmObject), the probeDisplayManager will query the probeLibrary for a ProbeMap. If there is not a *ProbeMap* installed in the probeLibrary for the class of that object, it will creat a *DefaultProbeMap*. It will then create a ProbeDisplay to the target object, and from then on, until the Probe Display is removed it will update the ProbeDisplay when receiving an update message.

Interface definition
(simtools/ProbeDisplayManager.html)
from simtools
(simtools.html)

library.

Marcus G. Daniels <mgd@santafe.edu>
(mailto:mgd@santafe.edu)

Last modified: 1998-01-08