

Berkeley Madonna User's Guide

Version 10.4

September 15, 2022

Frank Marcoline, Michael Grabe, Smita Nayak, Tim Zahnley, George Oster, Robert Macey

Berkeley Madonna, Inc.

1025 Peralta Ave

Albany, CA 94706

<http://www.berkeleymadonna.com>

Table of Contents

INTRODUCTION	5
CITING BERKELEY MADONNA	5
SYSTEM REQUIREMENTS	5
MODEL FORMAT	6
BERKELEY MADONNA GRAPHICAL WINDOWS	6
THE EQUATION WINDOW	6
THE FLOWCHART WINDOW	8
THE PARAMETER WINDOW	9
THE GRAPH WINDOW	9
Creating Multiple Graph Windows	9
Specifying Which Variables to View	10
Changing Axis Settings	11
Using Graph Buttons	11
Printing Graphs and Tables	13
Exporting Graphs and Tables	13
Oscilloscope Mode (<i>not available in version 10.4</i>)	13
Run Information	14
Readout	14
Initial Conditions	14
Fast-Fourier Transform	15
Nullclines	16
Histogram	16
THE DATASETS WINDOW	16
Importing Datasets	16
Vector Datasets	17
Matrix Datasets	17
Modifying and Removing Datasets	18
THE NOTES WINDOW	18
Embedding Pictures and Objects	18
RUNNING MODELS	19
SINGLE RUN	19
SLIDERS	19
BATCH RUNS	20
MONTE CARLO	20
PARAMETER PLOT	21
CHECK DT (<i>NOT AVAILABLE IN VERSION 10.4</i>)	23
FLOATING-POINT EXCEPTIONS	23
EQUATION SYNTAX	23
BASIC SYNTAX	23
DIFFERENTIAL EQUATIONS	25
DIFFERENCE EQUATIONS	26
DISCRETE EQUATIONS	27

Conveyors.....	27
Ovens	28
Queues	29
ROOT FINDER EQUATIONS.....	30
Initial and Step Root Finders	31
The Root Finder Algorithm.....	31
Multidimensional Root Finders	32
LIMIT EQUATIONS.....	32
OTHER EQUATION TYPES	33
METHOD Statement.....	33
DISPLAY Statement	33
RENAME Statement	33
BUILT-IN SYMBOLS.....	34
OPERATORS.....	34
BUILT-IN FUNCTIONS.....	35
ARRAYSUM, ARRAYMEAN, ARRAYSTDDEV.....	38
BINOMIAL, NORMAL, POISSON, RANDOM.....	38
DELAY	38
NETFLOW	38
PULSE	38
SQUAREPULSE.....	39
STEPSIZE	39
MAXINFLOW, OUTFLOW, OSTATE, QPOP, QPOPS	39
GRAPH	39
ARRAY EQUATIONS.....	40
DATASET FUNCTIONS.....	43
Plotting Matrix Datasets.....	44
TIPS AND TECHNIQUES.....	44
Creating Periodic Functions.....	44
Creating Piecewise-Linear Functions	44
Non-negative Reservoirs and Unidirectional Flows	45
INTEGRATION METHODS	45
FIXED-STEP SIZE METHODS.....	45
VARIABLE-STEP SIZE METHODS	45
USING DTOUT.....	47
CUSTOM DT METHOD	48
OTHER FEATURES	49
CHEMICAL REACTIONS	49
Adding Other Equations.....	50
CURVE FITTER.....	51
Multiple Fits.....	52
Specifying Initial Guesses.....	53
Floating-point Exceptions	53
OPTIMIZER	53
BOUNDARY VALUE SOLVER	54

SENSITIVITY	55
FLOWCHART REFERENCE.....	55
ALIASES.....	55
SUBMODELS.....	57
Creating and Deleting Submodels	57
Making Connections Between Submodels	58
Moving Icons Between Submodels.....	59
FLOWCHART TOOLBAR.....	60
FLOWCHART COMMANDS.....	63
FLOWCHART PREFERENCES	65
CUSTOMIZING ICONS.....	65
KNOWN BUGS AND LIMITATIONS.....	67

Introduction

Berkeley Madonna is a program that numerically solves systems of ordinary differential equations (ODEs), difference equations, and stochastic differential equations. The software lets users build mathematical models via a simple text editor or through an intuitive graphical user interface (GUI) that allows for the construction of complex models by assembling and linking together graphical symbols. As the visual model is created the corresponding equations are assembled automatically. Model solutions are then numerically generated using standard mathematical algorithms that have been optimized for speed. Once solutions are obtained, Berkeley Madonna provides a suite of graphical tools that allow users to easily represent their results and prepare them for incorporation into presentations and written documents.

Because of its user-friendliness, Berkeley Madonna is used at universities around the world to teach undergraduates, graduates, and professional students with little computational training how to construct and simulate models from the scientific literature. At the same time, its computational power and speed has made Berkeley Madonna an essential tool for academic scientists doing cutting edge research from wildlife conservation to engineering more efficient batteries. Additionally, Berkeley Madonna is used extensively in professional settings such as the pharmaceutical industry and biotech field.

Berkeley Madonna version 10 is a 64-bit program available for macOS and Windows operating systems. While these versions are similar, there are some differences, which are noted throughout this guide.

Citing Berkeley Madonna

When using Berkeley Madonna versions 9 or 10 in your work (publications, formal reports, class room materials, grant proposals, etc.) please cite the following article:

Marcoline, Furth, Nayak, Grabe, and Macey. (2022) CPT Pharmacometrics Systems Pharmacology. 11:290-301

This work can be found on the web for free at the following link:

<https://ascpt.onlinelibrary.wiley.com/doi/epdf/10.1002/psp4.12757>

System Requirements

macOS: Berkeley Madonna version 10 is tested on macOS 10.12 (Sierra) through 12 (Monterey), and may not work on earlier macOS operating systems. To run Madonna on earlier macOS versions, please contact us to obtain Berkeley Madonna version 9. Berkeley Madonna comes bundled with Java 17. Berkeley Madonna version 10 also requires Apple's Xcode command line tools, which is launched automatically when first run and downloaded from the Apple App Store. In order for this download to work properly, you must have an active internet connection the first time you run the software. This automatic installation of the Xcode tools is equivalent to running "`xcode-select --install`" in a Terminal.

Windows: PC compatible computer running Microsoft Windows 8 or later, including Windows 10. To run under Windows 7, please double click BerkeleyMadonna.bat instead of BerkeleyMadonna.exe. Berkeley Madonna comes bundled with Java 17 and Clang.

Memory: By default, Berkeley Madonna will allocate at most ½ of your total RAM. If this,

amount of memory is insufficient for your model, and you cannot reduce memory usage in the model through other means, which is always preferable, please contact us at *MadonnaExec 'at' gmail 'dot' com* for advice on increasing Berkeley Madonna's memory allocation.

Model Format

You can save your Berkeley Madonna model to a file (ending in “.mmd”) so that you can quit the software and come back to your work later or share your work with others. All of the information about your current session is stored in this model file including the equations, current parameter settings, visual models, and graphs. In version 9, we switched from an obscure binary format to an XML-like text format for transparency and easy future support and improvements. All Berkeley Madonna models created in versions 8 or earlier must be converted to the latest file format prior to running them in version 9 or 10. On Windows operating systems, conversion happens automatically upon loading the model, but this conversion is not possible on macOS. Therefore, if you are a macOS user, and you have old Berkeley Madonna models prior to version 9, you have two options for converting them:

1. Install Berkeley Madonna version 9 or 10 on a Windows machine, and then open the old model on that computer. Berkeley Madonna will then create a new converted file that can then be moved to the Mac and opened in version 9 or 10. *Please note, you do not have to have a valid Berkeley Madonna registration key for Windows to carry out this conversion – even unregistered copies have this conversion capability.*
2. Email your model(s) to *MadonnaExec 'at' gmail 'dot' com* with the subject title ‘PLEASE CONVERT’, and we will convert your models in 2-5 working days.

If neither of these two options work for your organization, please contact us at *MadonnaExec 'at' gmail 'dot' com*, and we can identify a solution with you.

Berkeley Madonna Graphical Windows

Berkeley Madonna uses various kinds of windows to represent your model including equation, flowchart, parameter, graph, datasets, and notes windows.

The Equation Window

You use the equation window to edit your model's equations. The editor is a simple plain text editor similar to SimpleText for Macs or Notepad for Windows. Once a model has been created, the equation window can be displayed at any time by choosing **Equations** from the **Model** menu.

When you create a new model using the **New Document** or **New Flowchart Document** command in the **File** menu, Berkeley Madonna creates a new, untitled equation window. You can only add equations to equations windows that do not have a graphical flowchart associated with them.

You need not keep the equation window open. As long as at least one other window for your model is open (parameters, flowchart, graphs, etc.), you can close the equation window and your model will remain open.

Changes made to your equations do not take effect until your model has been recompiled. There are several ways to recompile equations automatically by performing the following:

Model menu:

Compile Model

Run Model

Compute menu:

Compile Model

Check DT/TOLERANCE... (*not available in version 10.4*)

Run

Graph menu:

New Graph

Choose Variables...

Parameters menu:

Parameters

Parameter Sets

Define Sliders...

Show Sliders

Batch Runs...

Repeat Batch Runs

Monte Carlo...

Curve Fit...

Optimize...

Parameter Plot...

Boundary Value ODE...

Sensitivity...

However, if you unselect the **Automatically Recompile Equations** option in the **General** page of the **Preferences** dialog, Berkeley Madonna will only recompile your changed equations when you explicitly tell it to do so by choosing **Compile Model** from the **Model** or **Compute** menus.

By default, Berkeley Madonna will *not* prompt you before recompiling the equations, as it is generally very fast. However, if do not want the software to automatically recompile, select **Confirm Before Recompiling Equations** option in the **General** page of the **Preferences** dialog.

When Berkeley Madonna recompiles your equations, all runs in memory are discarded. In addition, any settings in your model that depend on particular symbol names (such as variables in graph windows, modified parameter values, sliders, etc.) will be lost if you changed the name of the symbol in your equations.

If Berkeley Madonna finds an error when compiling your equations, an error message is displayed and the suspect text will be selected in the equation window. The Equation Syntax section on page 23 explains how to write properly-formed equations.

When working with non-visual models, the equation window supports several editing features of note:

- Standard clipboard commands (cut / copy / paste) and drag-drop editing are provided.
- Text that extends past the right edge of the window is automatically wrapped to the next line; there is no need for a horizontal scroll bar and Berkeley Madonna doesn't provide one.
- The **Find...** and **Replace** commands in the **Edit** menu help you search for and replace text in your equations.
- The **Edit/Text Style**, **Edit/Text Size**, and **Edit/Text Font** commands allow you to change properties of the text used to display your model's equations. The text must be selected prior to using these commands either by using your mouse to highlight specific equations or by using the command **Edit/Select All** while the Equations Window is selected. By default, newly created models use a 12 point serif font for their equations.
- Select **Berkeley Madonna/Preferences...** and navigate to the **General tab** to change the **Equations Font Size**.
- To print the equation select **Print All** from the **File** menu allows you to print your model's equations or save them to a PDF. This command will cycle through your equations, graphs, and parameter window allowing you to print or save each one separately.
- The **Edit/Insert Image...** command allows you to embed pictures and objects within your equations. See Embedding Pictures and Objects on page 18. Note that embedded objects are treated by the compiler as if they were blanks (space characters). Therefore, they have no effect on the semantics of your model's equations.

As noted, many of these text editing features are disabled for visual models since Berkeley Madonna does not permit direct manipulation of the equations in these cases.

The Flowchart Window

The flowchart window provides an alternative to the equation window for constructing models. Instead of typing equations by hand, you build models by dragging icons from a toolbar onto a flowchart and connecting them with arcs to represent dependencies. As you construct your model graphically, Berkeley Madonna generates textual equations representing your model's structure.

To create a new visual model, choose **New Flowchart Document** from the **File** menu. This opens an empty flowchart window on which you construct your model. Unlike plain-text models (created with the **New Document** command), a visual model's equations cannot be edited directly in the equation window. Instead, equations are edited within each icon's "icon dialog" which is opened by double-clicking the icon.

The flowchart window can be hidden by clicking the **X** at the top of the window. If the window is hidden, choose **Show Flowchart** from the **Flowchart** menu.

Visual models can be converted into plain-text models by choosing **Discard Flowchart** from the **Flowchart** menu. After discarding the flowchart, the model's equations can be edited

directly in the equation window. Use caution as the flowchart cannot be recovered if you save a model after its flowchart has been discarded.

Berkeley Madonna cannot automatically convert plain-text models into visual models. However, you can create a new visual model that duplicates the behavior of an existing plain-text model, to do this, please refer to “Use the flowchart” in the *How Do I* in the help section to familiarize yourself with basic techniques for constructing and editing visual models. After you’re comfortable with the basics, refer to the Flowchart Reference on page 55 to learn about more sophisticated flowchart features.

The Parameter Window

You can display the parameter window by choosing **Parameters** from the **Parameters** menu. This window allows you to change parameters in your model as well as the integration method¹ without recompiling your model.


Parameters are symbols whose values do not depend on any other symbols in your model and do not change over time. Berkeley Madonna defines a set of built-in “system parameters” that exist in every model (STARTTIME, STOPTIME, DT, DTOUT, etc). These parameters always appear at the beginning of the list in the parameter window, and they vary slightly depending on the integration method.

To change the value of a parameter, select it with the mouse and type in a new value. The resulting value is displayed next to the parameter’s name in the list. If you enter an invalid expression, the value in the list won’t change.

Berkeley Madonna displays an asterisk (*) next to the name of each parameter that has been changed from its default value as specified in the equations. You can return a parameter to its default value by selecting it and clicking the **Reset** button.

The sliders window provides another convenient way to change parameters and automatically run your model. See Sliders on page 19 for details.

The Graph Window

When you run your model for the first time, Berkeley Madonna automatically creates a graph window. The window title indicates the run number and the variables plotted on the X and Y axes. If **Overlay Plots**  is off, each subsequent run will replace the data from the previous run. If it is on, subsequent runs will be added to the graph window. You can discard the most recent run by choosing **Discard Last Run** from the **Graph** menu. To discard all the runs shown in the window, choose **Discard All Runs** from the **Graph** menu.


Creating Multiple Graph Windows

Berkeley Madonna allows you to create more than one graph window for your model. This feature is useful if you want to view several variables with very different scales. To create another graph window, you have two choices. First, you can create an empty graph window

¹ The integration method can also be changed via the **Integration Method** submenu in the **Compute** menu.


by choosing **New Graph** from the **Graph** menu and add variables to it using the **Choose Variables** command. Note that new graph windows do not contain any run data and remain empty until you run your model again.

You can also create additional graph windows by selecting an existing graph window and choosing **Duplicate Window** from **Graph** menu. A graph window created in this way contains the same runs, variables, and other settings as the window from which it was duplicated.

When you run your model, Berkeley Madonna stores the results in all unlocked graph windows. If you want to preserve the contents of a graph window when you run your model, simply depress the **Lock**  button. Locked windows are never modified when you run your model.

For example, say you want to run your model twice and display the results of each run in a separate window. This can be accomplished by locking the graph window after performing the first run. Then, when you perform the second run, Berkeley Madonna creates a new graph window because no unlocked graph windows exist in which to store the new run.

The **Discard Last Run** command discards the last run shown in the active graph window. This run will also be discarded from all other unlocked graph windows in which it appears. The **Discard All Runs** command discards all runs shown in the active graph window. These runs will also be discarded from all other unlocked graph windows in which they appear. If a locked graph window is active and you attempt to discard runs, Berkeley Madonna asks you to confirm the operation.

Note that when using **Overlay Plots**  all graphed quantities from a single run are shown in the same color for clarity, and this color will likely differ from the color coding of the variable on the Y axis prior to overlaying.

Specifying Which Variables to View

To change which variables are plotted, select **Choose Variables...** from the **Graph** menu. You can specify any number of variables to plot on the Y axes. Each variable's scale and label can be shown on the left-hand or right-hand Y axis. To change which axis is used, select the variable and toggle the **Right Axis** check box. To hide or show a variable, select it and toggle the **Visible** check box.

By default, Berkeley Madonna plots TIME on the X axis. You can change this to any scalar variable in your model using the X axis control in the **Choose Variables...** dialog. If you are working with arrays, you can plot their final values (i.e., when TIME = STOPTIME) against the element index by choosing **[i]** in the X axis control.

When you specify new variables for your plot, they may not appear after you click OK in the **Choose Variables...** dialog. This happens because Berkeley Madonna only stores the variables in your model that appear in the Y Axes list for at least one graph window at the time your model is run. If you add new variables after the run, you will need to *run your model again* so they are stored in memory. For example, say your model has three variables: A, B, and C. Assume that the first time you run your model, you specify variable A on the Y axis. Berkeley Madonna stores values of A for this run, but not the values of B or C. Then, if you change the graph to plot B and C in addition to A, only a curve for A will appear. If you run your model again, A, B, and C will be stored and will appear in the plot.


If you find yourself working with a group of variables but do not want to view all of them at once, use the variable buttons at the bottom of the graph window (available when the **Variable Buttons** command in the **Graph** menu is checked). Clicking a variable button shows or hides the variable. Holding the shift key while clicking a variable button moves the variable from one Y axis to the other.

Changing Axis Settings

By default, Berkeley Madonna automatically chooses the scale for each axis in your plot. You can override automatic scaling by choosing **Axis Settings...** from the **Graph** menu and selecting the **Scales** tab.² When one of the **auto** boxes is checked, the settings for that axis are grayed because Berkeley Madonna is choosing them for you. To change them, you must first uncheck the **auto** box. The **# Div** field specifies the number of divisions the axis is divided into by the grid lines.

Checking a **log** box in the **Scales** page of the **Axis Settings...** dialog changes the scaling for that axis from linear to logarithmic. If the **auto** box is checked, Berkeley Madonna will choose the axis range for you using exact powers of ten. If you uncheck the **auto** box, you can specify whatever limits you want. Note that when using logarithmic scales, Berkeley Madonna does not distinguish between positive and negative values. Also, zero values are not plotted at all since the logarithm of zero is infinite.

You can quickly “zoom in” to view a portion of your data. To do this, click and drag the mouse to draw a rectangle over the portion you wish to view (rectangle may not appear on some versions of macOS). When you release the mouse, Berkeley Madonna adjusts the axis limits so that the portion that was in the rectangle now fills the window.

To undo the effect of a zoom-in operation, click the **Zoom Out**  button. Each click of this button while holding down the command key [macOS] or Alt key [Windows] undoes the effect of the most recent zoom-in operation. Simply clicking the **Zoom Out** button alone undoes the effect of *all* zoom-in operations.

Any changes made in the **Scales** page of the **Axis Settings...** dialog while zoomed-in are lost when you zoom out. To prevent this from happening, zoom out all the way before making changes to the axis scales.

By default, Berkeley Madonna chooses the text used to label the axes based on the variables you are plotting. You can override the default by choosing **Labels** tab in the **Axis Settings...** dialog.³ Uncheck the **default** box for each label you want to change and edit its text.

Using Graph Buttons

The buttons in the top-left corner of the graph window perform the following functions:

² You can open the Scales page of the Axis Settings... dialog in one step by double-clicking any axis in the graph window.

³ You can open the Labels page of the Axis Settings dialog in one step by double-clicking any axis label in the graph window.



Run. Runs the model once. Same as choosing **Run** from the **Model** or **Compute** menus, except for parameter plot windows where this button performs a parameter plot run.



Lock. Prevents the window's contents from being changed when running the model or discarding runs in other windows.



Overlay Plots. When enabled, new runs are added to existing runs in the window. Otherwise, existing runs are discarded before new runs are stored.



Table. Displays the contents of the window in tabular form. Note that the TIME column reflects the time for the run shown in the rightmost column of the table. If there are other runs with different time scales, their values will be mapped to the times shown in the TIME column using linear interpolation.



FFT. Displays the contents of the graph window in the frequency domain. See Fast-Fourier Transform on page 15.



Legend. Displays the legend. The legend shows the variable name and run number for each curve. Additional information in parentheses is included for Batch Runs: the parameter value for separate runs and labels for mean and mean±sd runs. You can place the legend anywhere in the graph window by dragging it with the mouse.



Parameters. Displays the parameter list. This list shows the value of each parameter when the model ran. If there is more than one run in memory and a parameter is not the same in all runs, the minimum and maximum values of the parameter are shown. You can position the parameter list with the mouse just like the legend.

NA **Oscilloscope Mode.** (*not available in version 10.4*)



Line Appearance. Displays curves using thin lines, thick lines, or data points. Also use this icon to choose a color palate: Default, Bright, Bold, Accessible, or Gradient.



Grid. Shows the grid lines.



Readout. Displays X and Y values as you drag the mouse over the plot, instead of zooming in. See Readout on page 14.




Initial Conditions. Enables you to set initial conditions by dragging the mouse over the plot. See Initial Conditions on page 14.



Zoom Out. Undoes the effect of the last zoom-in operation. If no zoom-in operations have been performed, this button is grayed.



Note. Create text on the graph.



Nullclines. Plot the nullclines of your system of ODEs. These solutions can only be rendered as data points. See Nullclines on page 16.



Histogram. Create a histogram of a time dependent variable. This button opens up a new graph window and automatically bins all time dependent variables. The bin number can be automatically adjusted. Only time dependent variables can be graphed in this manner. See Histogram on page 16.

Details of the graph window can be modified by accessing the **Graph** page of the **Preferences...** dialog. For instance, you can:

1. Choose to plot all Imported Data with Circles
2. Adjust the axis label font sizes
3. Make the Background Gray (does not impact current graph(s))

Once modified, these features will always be present when new graphs are created.

Printing Graphs and Tables

You can print a graph or table by selecting its window and choosing **Print Graph...** or **Print Table...** from the **File** menu. Before printing, however, you can choose **Print Preview Table/Graph...** to see what the image will look like. From this window you can change from portrait to landscape and scale the size of the image. Additionally, you can choose **Page Setup...** from the **File** menu to manipulate the image output.

On Windows, you control the position of the output on the printed page by specifying margins in the **Page Setup** dialog. On macOS, select **Paper Size:Manage Custom Sizes...** to specify margins.

Exporting Graphs and Tables

You can save graphs as pictures by choosing **Save Graph...** from the **File** menu. Note that you must first select the graph window you want to save. Graphs are saved as jpeg images.

Tables can be saved as tab-delimited text (*.txt) or comma separated values (*.csv) files by activating the window containing the table and choosing **Save Table...** from the **File** menu. Tables can also be copied to the clipboard by highlighting the text manually, then selecting the **Copy** command from the **Edit** menu. The information can then be ported to another document with the standard paste command. Tables are always copied in tab-delimited text format.

Oscilloscope Mode *(not available in version 10.4)*

Oscilloscope Mode causes the graph to start drawing at zero on the X axis whenever a trigger event occurs. To use this feature, define a variable in your model named "trigger". Berkeley Madonna recognizes a trigger event whenever this variable changes from zero to any nonzero

value. For example, to trigger whenever the voltage V_m is greater than 15 millivolts, add this equation to your model:

$$\text{trigger} = V_m > .015$$

Compile and run your model. Then, click the **Oscilloscope Mode** button. The variable(s) you chose will be displayed as one or more curves starting at time zero depending on how many trigger events occurred. To restore normal (non-oscilloscope) mode, click the **Oscilloscope Mode** button again.

Note that oscilloscope mode is disabled when the Fast Fourier Transform is used.


Run Information

By default, Berkeley Madonna displays the number of steps and elapsed time of the last run in the lower left-hand corner of the graph window. The number of steps reflects how many times your model was stepped forward in time after initialization. This number is one less than the number of data points shown in the graph or table since it does not include the initialization step. However, if you use *DTOUT* to reduce the number of steps stored in memory, the data points presented on the graph and in the table will be *less* than the number shown in the lower left corner. This is because the graphed output has been subsampled after the run to reduce memory loads. See Using *DTOUT* on page 47.

When performing batch runs with averaging, the run information indicates the total number of steps and elapsed time for *all* of the individual runs used to create the resulting average. The same totaling also applies to parameter plot runs.

Run information is not shown when either the Readout or Initial Conditions features are active.


Readout

Berkeley Madonna can display the coordinates of any point in the graph window. Activate readout mode by depressing the **Readout**  button, then move the mouse to an area of interest on the graph. Click the mouse button and the corresponding TIME, Left Y Axis, and Right Y Axis values are displayed at the bottom left of the window.

When Fast Fourier Transform mode is active, the readout displays both the frequency and period corresponding to the X axis coordinate. Additionally, the Left and Right Y Axis values are also given.

Note that you must deactivate Readout mode to use the “zoom in” feature.

Initial Conditions

The **Initial Conditions**  button provides a convenient way to set the initial values of two reservoirs or difference equations in your model by dragging the mouse over the X-Y phase plane. To use this feature, you must:


- Set the X and Y axis variables to reservoirs or difference equations in your model. The initial value expressions for these variables must be parameters that can be accessed from the parameter window. For example, if you plot reservoir “A” on the X axis and reservoir

“B” on the Y axis, both “INIT A = ...” and “INIT B = ...” must appear in the parameter window.

- Manually set the scale for both the X and Y axes using the **Axis Settings** dialog.

Once you are plotting the proper variables on each axis and have defined the scales, the **Initial Conditions** button will be enabled. Click it and the readout cursor will appear. Move the cursor to a point in the phase plane and click the mouse button. Berkeley Madonna will set the initial values of the X and Y axis variables and run your model.


Below is an example with the U and V nullclines shown as blue and green and the red phase plane trajectory going to a fixed point at $X = 0.25$, $Y = 4.75$.



Phase plane portrait created using the nullcline feature together with initial conditions.

Fast-Fourier Transform

When investigating oscillating systems, it is useful to be able to find the period or frequency of a periodic solution. One way to do this is by looking at the Fourier transform of a time series run, which gives the contribution to the data from each periodic component.


To view the Fourier transform of your data, click the **Fast Fourier Transform**  button in the graph window. The data will be transformed into the frequency domain. To find the period and frequency, activate readout mode by clicking the **Readout** button, place the mouse pointer over the lowest frequency peak, and press the mouse button. The frequency, period, Left Y Axis value, and Right Y Axis value will be displayed in the lower-left corner of the window.

Once you have transformed your data, Berkeley Madonna keeps the frequency data in memory so you can quickly switch between the time and frequency domains. It also keeps track of the axis scales used in the time and frequency domains separately so you do not have to change the scale every time you switch between domains.


When a run is saved in a graph window where **Fast Fourier Transform** is active, the run will be transformed automatically as it is saved. If you do not want this to happen, switch back to the time domain or lock the window.

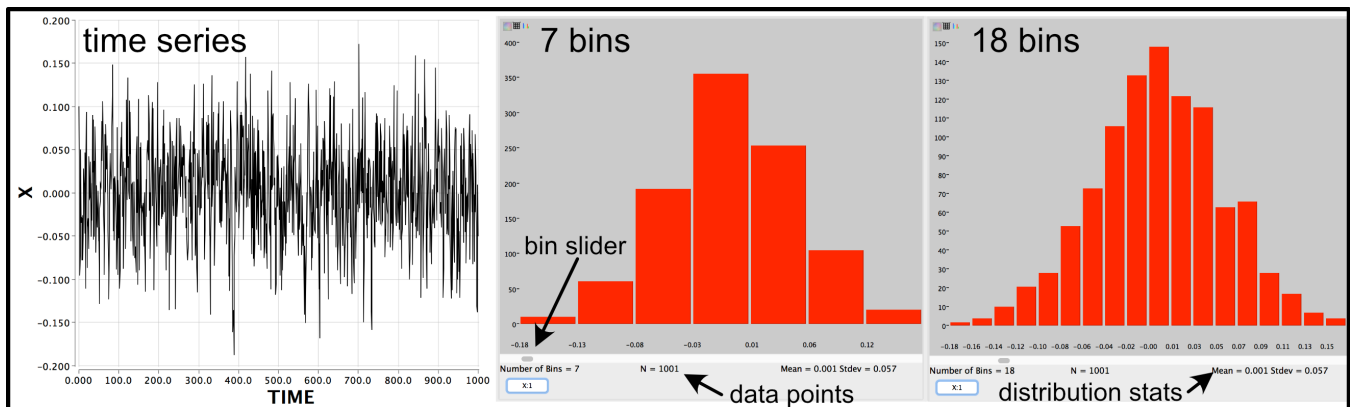
Nullclines

Given the ODE $x' = f(x,y)$, the nullcline corresponding to x is the geometric shape in the x,y -plane for which $f(x,y) = 0$. For a set of ODEs with a second equation $y' = g(x,y)$, the intersection of the x -nullcline and y -nullcline ($g(x,y) = 0$) corresponds to the equilibrium points of the system. To visualize the nullclines in this example, plot y on the y -axis and x on the x -axis (double click anywhere in the graph and replace TIME with x on the X-Axis). Now press

$y=0$ and the nullclines for g and $f = 0$ will be computed and plotted using dots (each curve will have a different color, and select **Legend** to identify which nullcline is which. Finally, use the Initial Conditions tool  to determine how the solution behaves with respect to the nullclines. For hands-on experience, open the “How Do I Plot Nullclines” model in the Help menu.

Histogram

From version 9 on, Berkeley Madonna allows the user to quickly and easily plot time series data in histogram format by pressing the Histogram button  from the graphical toolbar. The data points will be binned between the maximum and minimum values, and the number of bins can be changed by sliding the bin bar at the bottom of the histogram as shown in the figure. Only a single variable can be shown at once, and use the toggle bottom at the bottom of the window to change what is displayed.



The Datasets Window

You can import numerical data from a text file into your model. The datasets window lists the datasets you have already imported. Datasets can be used as targets for curve fitting (see **Curve Fitter** on page 51) and as piecewise-linear functions in your model's equations (see **Data Set Functions** on page 43). These data sets must be ASCII formatted files.

Importing Datasets

To import a dataset into your model, choose **Import Dataset...** from the **File** menu. Or, open the datasets window by choosing **Datasets** from the **Model** menu and click the **Import...** button. Then, choose a text file containing the numerical data you want to import. The data must be in tab-delimited text format or comma-separated values (*.csv) format. Data files in

these formats can be generated by spreadsheet programs such as Microsoft Excel by specifying the appropriate file type in the “Save As” dialog.


Once you have selected the data file, the **Import Datasets** dialog window appears where you specify the name by which the dataset is referred to in your model and the type of dataset (Vector (1D) or Matrix (2D)) to create.

Berkeley Madonna assumes that the data file consists of a rectangular array of numerical values (integers and/or decimal numbers). Rows are read sequentially starting at the beginning of the file. Values in each row are read left to right until either the end of the row is reached or a non-numerical value (a letter or symbol, for example) is read. Berkeley Madonna then checks that this row contains the same number of numerical values as the previous rows. If it does not, you will get an error message stating “number of columns is not constant”. Otherwise, Berkeley Madonna continues reading rows until the end of the file is reached. Note that rows containing no numerical values are ignored. This makes it possible to have a line of text in your data file (such as a title at the beginning or a text break at certain intervals throughout the file).

Once the data has been entered, if it is 1D, it will automatically be graphed in a Graph window as circles or data points. You can change this default representation in the **Graph** window of **Preferences...** under **Berkeley Madonna** in the file menu. Later on, you can always add or remove the data points to a particular graph using **Choose Variables...** from the **Graph** menu item.

Vector Datasets

Vector (1D) datasets (also known as one-dimensional datasets) map an input value (X) into an output value (Y). Two columns of data are required. When creating a vector dataset, you specify which columns in your imported data file to use for the input domain (X column) and output range (Y column). The X column must consist of monotonically increasing values (i.e., the value in row n+1 must be greater than the value in row n). If not you will get the following error:



The datasets window lists each vector dataset with its name followed by the number of points in parentheses.

Matrix Datasets

Matrix (2D) datasets (also known as two-dimensional datasets) map two input values (X and Y) into an output value (Z). Berkeley Madonna uses your *entire* imported data file to construct the dataset using the following scheme. For example, assume that you have imported a data file with 35 elements in five rows and seven columns:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} & a_{17} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} & a_{27} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} & a_{37} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} & a_{47} \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} & a_{56} & a_{57} \end{bmatrix}$$

Berkeley Madonna uses elements in the first column (a_{21} through a_{51}) to define the dataset's X domain and the elements in the first row (a_{12} through a_{17}) to define the dataset's Y domain.

The remaining elements (a_{22} through a_{57}) define the dataset's output range (Z) for the corresponding X and Y input values. Note that although element a_{11} is not used, it must still be present in your data file or Berkeley Madonna will refuse to import the file.

The datasets window lists each matrix dataset with its name followed by the dimensions of the dataset in square brackets.

Note that matrix datasets cannot be plotted directly in graph windows. This is because Berkeley Madonna does not have a 3D plotting capability. However, you can still get an idea of what a matrix dataset looks like using the technique described in Plotting Matrix Datasets on page 44.

Modifying and Removing Datasets

A dataset's name can be changed by selecting it in the datasets window and clicking the **Rename...** button.

A dataset can be removed from your model by selecting it in the datasets window and clicking the **Discard...** button. You may want to remove datasets that you are no longer using since they can significantly increase the size of your model file depending on the number of elements.

To modify the contents of a dataset, simply import a new dataset and give it the same name as the existing dataset. Berkeley Madonna will confirm that you want to replace the contents of the dataset.

The Notes Window

The notes window lets you save text and graphic notes in your model without cluttering your equations with comments. Notes can be displayed at any time by choosing **Notes** from the **Edit** or **Model** menus. Like the equation window, the notes window supports drag & drop text editing, find/replace, embedded pictures and objects, etc. You can also print your notes or save them as a text file.

Unlike the equation window, you can mix multiple styles of text within your notes by using the font name, size, and style controls at the top of the notes window.


Embedding Pictures and Objects


Images can be embedded in your notes and equations windows. Choose the **Insert Image...** command in the **Edit** menu to accomplish this task, and navigate to the location of your image file on your computer. To insert in the equations window, first click on that window prior to

selecting **Edit/Insert Image...**, and similarly, to insert in the notes window click on that window first. After insertion, manipulation of the images is limited.

Running Models

Single Run

The simplest way to run your model is to click the **Run** button  in the flow chart, parameter, or graph windows. Or you can choose **Run Model** from the **Model** menu item or **Run** from the **Compute** menu item. The result of the run is displayed in a graph window. If no graph exists, Berkeley Madonna will create one and plot the first eight variables⁴ in your model.

Each time you run your model, the data from the new run will replace the data from the previous run. If you want to display more than one run at a time, depress the **Overlay Plots**  button.

When **Overlay Plots** is on, each new run is added to the previous runs in memory. The title of the graph window indicates how many runs are in memory.

When **Overlay Plots** is off, Berkeley Madonna discards all existing runs in memory when a new run is completed.

Sliders

Sliders provide a quick way to change a parameter and run the model in one step. To create sliders, choose **Define Sliders...** from the **Parameters** menu. In the dialog, select a parameter and add it to the sliders list by clicking **Add**. Select linear or logarithmic scaling and adjust the minimum, maximum, and increment or multiplier for the slider. If you want additional sliders, add them now. Click **OK**. Berkeley Madonna will display the sliders you defined at the bottom of all current graph windows.

To change a parameter via a slider, drag the indicator or click the arrows. When the mouse button is released, the parameter change will be applied and the model will rerun.

The value range of an individual slider can be quickly changed by double-clicking anywhere on the slider's text (where the parameter name and value are displayed at the bottom of the graph window). This action opens the **Define Sliders...** dialog and selects the slider you clicked.

Each slider has a check box labeled **10x** which reduces the range of the slider by a factor of ten around its current value. This feature makes it easier to zero in on a desired value. Note that this 10x mode is cancelled if you make any changes via the **Define Sliders** dialog or recompile your model.

⁴ Only the first two variables will be visible. To show the other variables, use the variable buttons at the bottom of the graph window.

You can hide the sliders by choosing **Hide Sliders** from the **Parameters** menu. You can bring the previously-defined sliders back by choosing **Show Sliders** from the **Parameters** menu.

Batch Runs

Berkeley Madonna can automatically run your model multiple times while optionally varying a specified parameter. To use this feature, choose **Batch Runs...** from the **Parameters** menu. Berkeley Madonna displays the **Batch Runs...** dialog. Set it up as follows:

1. Specify the parameter you want to vary for each run from the **Parameter** control. If you do not want Berkeley Madonna to change a parameter, choose None. None is convenient for running many stochastic models that show variability even when all parameters remain the same.
2. Specify the number of runs you want to perform. If you choose to plot each run with **Keep Runs Separate**, you may run into graphing problems with too many data points. If this arises, reduce the number of runs or change your DTOUT so that you plot fewer points per trace.
3. Specify the values you want the parameter to take for the first and last run. Enter these values in the **Initial Value** and **Final Value** fields, respectively. Note that these values must be different.
4. Specify whether the parameter should be varied according to an arithmetic or geometric series by clicking one of the **Series Type** radio buttons. Note that the initial and final values for a geometric series must be nonzero and the same sign.
5. Specify how Berkeley Madonna should process the runs. There are three choices:
 - **Keep Runs Separate.** Berkeley Madonna stores each run separately in memory and plots them on all unlocked graph windows as they complete, even if Overlay Plots is not selected on that graph.
 - **Compute Mean.** Berkeley Madonna computes the mean value at each time point of all runs and plots the result as a single run.
 - **Compute Mean \pm SD.** Berkeley Madonna computes the mean and standard deviation at each time point of all runs and stores the result as three runs: the first run is the mean, the second is the mean minus the standard deviation, and the third is the mean plus the standard deviation. Thus, the second and third runs show the bounds of the area extending one standard deviation on each side of the mean.
6. Click OK to start the runs.

Note that Berkeley Madonna will *add* the resulting run(s) to the graph if **Overlay Plots** is selected. If you want any previous runs to be discarded, be sure to turn off **Overlay Plots** before starting the runs.

You can easily repeat the previous Batch Runs command by choosing **Repeat Batch Runs** from the **Parameters** menu.

Monte Carlo

You can simulate stochastic elements of your models by using the Monte Carlo feature found in the Parameters drop down menu. Note that this feature is only available for fixed timestep

integration methods. The menu option is unavailable when the current integration method is either auto-stepsize or Rosenbrock (stiff). Selecting the Monte Carlo option will bring up a dialog box that allows you to turn parameters of interest into random variables that are assigned at the beginning of each run according to one of three different distributions (uniform, normal, log normal). Start by adding parameters from the left side of the box to the Selected section. For each selected parameter, assign it a distribution from the drop-down list on the far right of the dialog box, and then fill in the distribution details. Select the number of times to repeat the simulation by entering a value in the "Runs:" field. At the start of each new run, the selected parameters will be given a new value from their respective distributions. Next, select one of three run modes: keeps runs separate, compute mean, or compute mean \pm standard deviation (SD). When the computations complete a graph will appear displaying the results corresponding to the mode that was chosen. Selecting the histogram button will produce a second graph showing the distribution of the simulated random variables.

For each run, the selected parameters will be randomly assigned a value from the distribution selected on the far right of the dialog box (uniform, normal, log normal). For distributions other than the uniform distribution, you can additionally place artificial limits on the minimum and maximum value of a parameter. These limits are enforced by the rejection method: random variates of the distribution are repeatedly generated until the value falls within the specified minimum and maximum. This does not change the shape of the distribution within the selected domain (but it does change the normalization). This is often preferable to using a random number generator directly in your model equations and using LIMIT statements. When you constrain a variable with LIMIT statements, whenever it would fall outside the imposed limit, it is instead clamped to that limit, which distorts the shape of the distribution within the domain specified by the LIMIT statements.

The Keeps Runs Separate option will plot all repeated trials on the same graph over top of each other. Be careful with this when repeating many runs, as it may overwhelm the graph with data. The next two options compute the mean values of all model values for the entire time series, and then report this value. The \pm standard deviation option will plot 2 more curves over this mean value curve corresponding to the standard deviation of the value being plotted determined from taking all aggregate data for that parameter at each time point and computing the SD, which is then added/subtracted from the mean value.

There are few details that you should be aware of regarding these plots. First, when a mean curve is computed from repeated runs of a given process using different stochastically chosen rate constants, the underlying nature/shape of the curve likely changes. For instance, if you were to simulate many times the equation $dy/dt = -ky$ with k drawn from a uniform distribution from 1 to 5, each individual run would result in a pure exponential. However, averaging all runs together, as described above, results in curve that is no longer exponential in time. Second, consider the exponential decay of this system with $y(0) = 1$. The quantity $y(t)$ remains positive definite for all values of k (since k has values drawn from 1 to 5); however, $y - SD$ will take on negative values. In cases like this, you may want to represent your Monte Carlo results by keeping runs separate and showing the resulting spread. In future releases, we will add a confidence interval option that will more appropriately treat the distribution values in these cases.

Parameter Plot

The Parameter Plot feature enables you to plot the result of each run as a single data point over a range of parameter values. With this function, Berkeley Madonna provides a variety of methods to convert a time dependent variable's value over time into a single value that can be


plotted as a function of the parameter sweep. The initial, final, minimum, maximum, mean, and standard deviation calculations are self-explanatory. The oscillation frequency and amplitude calculations consist of performing an FFT on the run data and selecting the frequency component with the maximum amplitude.

To perform a parameter plot, choose **Parameter Plot...** from the **Parameters** menu. This opens the **Parameter Plot...** dialog where you specify the parameter to vary and the variables to plot. As part of selecting the parameter to vary, you specify the number of runs (steps), initial value, final value, and series type. This setup is identical to that used in the **Batch Runs...** dialog.

The variables portion of the **Parameter Plot...** dialog is similar to the **Choose Variables...** dialog that you are likely familiar with from adding elements to graphs: you select the variables you want to plot and add them to the **Y Axes** list. However, for each variable, you also must select at least one of type of output conversion to generate a single value from the run data. Simply check the appropriate box (or boxes) to the right of the **Y Axes** list. As each box is checked, the variable name in the list reflects the type of conversion (for example, 'min' for minimum).

Once you have specified the parameter to vary and the variables to plot, click the **Run** button in the dialog. Berkeley Madonna performs the series of runs and plots the result in a new graph window.

Since the independent variable in a parameter plot is a parameter in your model rather than **TIME**, Berkeley Madonna displays the result of the run in a special graph window known as a parameter plot graph window. Such windows can be identified by the text "**Parameter Plot...**" in the title bar. The x-axis label must be read to determine which parameter value was swept for a particular run. Parameter plot graph windows are different from normal graph windows in other ways, as well:

- They can only display results of parameter plot operations based on the varied parameter.
- The X axis variable cannot be changed.
- Since the window's contents are not based on **TIME**, you cannot perform a Fast Fourier Transform on them or use them in other features.
- The **Run**  button in the Graph Window (not in the Parameter Window) performs the last parameter plot operation as specified in the Parameter Plots dialog for this parameter. This makes it easy to perform the Parameter Plots operation again (for example, after changing some other parameter in the parameters window) without invoking the **Parameter Plot...** dialog box.

Once you have a parameter plot graph window, you can easily change the parameter's range and/or variables plotted by invoking the **Parameter Plot...** dialog again. This can be done in several ways:

- Double-click anywhere in the plot area of the graph (the same action that displays the **Choose Variables...** dialog in normal graph windows).
- Select the parameter plot graph window and choose either **Choose Variables...** from the **Graph** menu or **Parameter Plot...** from the **Parameters** menu.

When the **Parameter Plot...** dialog is opened for an existing parameter plot graph window, you will note an unchecked **Create New Window** button in the lower-center. This means that

any changes you make to the list of variables will be applied to the graph window for which the dialog was invoked. If you would rather not modify this window, check this box and Berkeley Madonna will create a new parameter plot graph window.

Check DT *(not available in version 10.4)*

The **Check DT/TOLERANCE** command in the **Compute** menu helps you verify that you are using a sufficiently small value of DT (when using fixed-stepsize integration methods) or TOLERANCE (when using variable-stepsize integration methods) for accurate results. Berkeley Madonna performs this check by running your model twice: first with the current value of DT or TOLERANCE and second with DT or TOLERANCE reduced by a factor ten. After both runs have been completed, Berkeley Madonna displays the **Check DT Report** or **Check TOLERANCE Report** dialog. These dialogs show the maximum and root-mean-square deviation between the two runs for each stored variable. Both runs are also stored in all unlocked graph windows so they can be compared visually.

Floating-Point Exceptions

If you get a floating-point exception when running your model, it means some numerical operation resulted in an error. Exceptions include division by zero, numerical overflow, taking the square root of a negative number, etc.

When the **Stop On Exception** box in the **Settings...** dialog (**Model** menu) is checked, which is the default, Berkeley Madonna checks for exceptions after each time step has been calculated, and its results are stored in memory. If an exception has occurred, the run is stopped and the exception is reported.

When the **Stop On Exception** box is not checked, Berkeley Madonna checks for exceptions after the final step has been completed (STOPTIME has been reached). In this case, it may not be possible to determine exactly when the first exception occurred.

Non-finite values are not displayed in graphs, so to find out where things started going wrong it may help to view the results in table form. Unlike graphs, tables display invalid numerical values.

Equation Syntax

To display a summary of Berkeley Madonna's equation syntax, choose **Equation Help** from the **Help** menu. You can select text in this window and drag it to your equation window (as long as you don't have an active flowchart window).

Basic Syntax

Berkeley Madonna models consist of a list of equations. Equations can span more than one line, and you can put multiple equations on a single line. The order in which equations appear is generally unimportant. Most equations are of the form:

$$\text{variable} = \text{expression}$$

Variable names consist of one or more characters. You may use letters (A-Z), digits (0-9), and underscore (`_`) in variable names. However, variable names must not begin with a digit.

A variable can be given almost any name as long as it doesn't clash with any of Berkeley Madonna's built-in function or symbol names (see Built-in Functions on page 35 and for visual models, use the RENAME command from the Globals Window).

Comparisons between variable names are case-insensitive. For example, the names "Result", "RESULT", and "result" all refer to the same variable.

Expressions use standard infix notation. That is, arithmetic and relational operators are placed between their two operands like this:

$$\text{result} = a + b * c$$

The usual operator precedence rules apply. For example, in the above equation, b and c are multiplied first, then their product is added to a. Parentheses can be used to force different order of evaluation. For example:

$$\text{result} = (a + b) * c$$

Berkeley Madonna provides a number of built-in functions. The name of the built-in is followed by one or more arguments enclosed in parentheses. However, functions which take no arguments (e.g., PI) are not followed by parentheses. For example:

$$\text{result} = a * \text{SIN}(2 * \text{PI} * x)$$

Multiple function arguments are separated by commas like this:

$$\text{result} = \text{SUM}(a, b, c) / n$$

You can place comments anywhere in the equations by enclosing them in curly brackets. For example:

```
{This is a sine wave}
result = a * SIN(x)
{This is the amplitude}
a = 7.5
```

Curly-bracket comments can span multiple lines and can be nested, like this:

```
a = ...    {define a}
{the following two lines have been disabled by this comment:
b = ...    {define b}
c = ...    {define c}}
d = ...
```

Berkeley Madonna also recognizes single-line comments beginning with the semicolon character. This type of comment extends through the end of the line and doesn't require an "end of comment" character like curly brackets do. For example:

```
a = ...    ;define a
b = ...    ;define b
c = ...    ;define c
```


Differential Equations

Equations like those shown above simply assign the value of the expression on the right-hand side to the variable on the left-hand side. Such equations are referred to as “formulas” by the flowchart editor.

Ordinary differential equations (known as “reservoirs” in the flowchart editor) are defined by two equations: an initializer equation and an integrator equation. The initializer equation determines the initial value of the reservoir. The integrator equation determines how much the reservoir’s value increases or decreases during each time step (also known as the inflow or outflow).

Berkeley Madonna supports different forms of initializer and integrator equations. The following three forms of initializer syntax are functionally identical; each initializes reservoir R to an initial value denoted by “...”:

```
R(STARTTIME) = ...
INIT R = ...
INIT (R) = ...
```

The following five forms of integrator syntax are also functionally identical; each defines a net flow into reservoir R denoted by “...”:

```
d/dt(R) = ...
R' = ...
FLOW R = ...
R(t) = R(t - dt) + (...) * dt
R = R + dt * (...)
```

We don’t recommend the use of the last two forms since the notation is more error-prone.

Since the different forms of these equations are equivalent, you can use whichever you like in your own equations.

The second integrator syntax shown above ($R' = \dots$) deserves special mention. Unlike other notations, this “prime” notation allows you to define higher-order systems without explicitly writing their equations. For example, you can define a third-order system like this:

$$u''' = 2*u'' - 5*u' + u - 1$$

Berkeley Madonna internally translates this into a system of three first-order equations like this:

$$\begin{aligned} u''' &= 2*u'' - 5*u' + u - 1 \\ d/dt(u'') &= u''' \\ d/dt(u') &= u'' \\ d/dt(u) &= u' \end{aligned}$$

The generated equations, although not shown in the equation window, do exist behind the scenes. This means that you have to provide initializers, like this:

```
INIT u'' = ...
INIT u' = ...
INIT u = ...
```

For instance, the solution to the following harmonic equation will produce a sine wave with amplitude 1:

$$x'' = -x$$

$$\text{INIT } x = 0$$

$$\text{INIT } x' = 1$$

Difference Equations

Berkeley Madonna also supports difference equations. Each difference equation consists of two equations: an initializer equation and a “next” equation. The initializer equation determines the initial value of the difference equation. The next equation determines the value the difference equation will take at the next time step. The syntax of the initializer equation is the same as used by reservoirs (see above). When typing equations into the equation window, there are two equivalent forms of the next equation that can be used:

$$v(t + dt) = \text{expression}$$

$$\text{NEXT } v = \text{expression}$$

where $\text{INIT } v = X$ must also be supplied as stated, and X is an initial value. The method of integration is important. For standard usage, a fixed time step method such as Euler, RK2, or RK4 will produce the same results, while Auto-stepsize or Rosenbrock methods will produce different results, and are unlikely to be useful. When Berkeley Madonna computes the next value of a difference equation, it uses the *current* values of any variables that appear in the expression on the right-hand side of the next equation. For instance, if the current value of v is 10.2, and the difference equation is:


$$v(t+dt) = v + 1,$$

then the next value of v will be 11.2 followed by 12.2, and so on. This progression is independent of the step size.

Difference next equations can be constructed from the Flowchart as well. First, construct a new flowchart, then select **Difference Equations** from the **Flowchart** menu. As you construct a model, the difference equations will be written. They take the following form in version 8 and version 9.1.16 and later (this is handled differently in version 9 prior to 9.1.16), and in the Equations Window you will see:

$$d/dt R = -J \quad \rightarrow \quad R(t+dt) = R - J,$$

for reservoir R and a flow J out of R , as in the following model:



where J is the flow out, and it can take on any mathematical expression. DT will automatically be set to 1, and the method will be Euler.

Discrete Equations

Berkeley Madonna supports three kinds of discrete objects: conveyors, ovens, and queues. These objects accumulate inputs, hold them for a period of time, and release them. The manner in which this is done differs for each kind of object.

Note that discrete objects should only be used with fixed-stepsize integration methods. Furthermore, since the behavior of these objects is discrete rather than continuous, no additional accuracy is gained by using Runge-Kutta integration methods. Therefore, models which employ these objects should use Euler's method.

Conveyors

Conveyors work like a conveyor belt. Each time the model is stepped, an input is placed on the conveyor along with a transit time. The transit time determines how long this input remains in the conveyor. When its transit time has expired, the input is released from the conveyor.

Conveyors are defined using the following syntax:

```
name = CONVEYOR(input, tt, cap)
```

where *name* is the name of the conveyor, *input* is the input expression, *tt* is the transit time, and *cap* is the capacity. The capacity argument is optional; if omitted, the conveyor has unlimited capacity; however, this feature only works when coupled with MAXINFLOW, as discussed below. The name of the conveyor can be followed by array dimensions in square brackets to create an array of conveyors.

The name of the conveyor is a variable which can be used in other equations or plotted. Its value is the sum of all inputs in the conveyor multiplied by DT.

To determine when inputs have exited the conveyor, use the OUTFLOW built-in:

```
output = OUTFLOW(name)
```

Since transit times for each input can vary, it is possible that more than one input will exit the conveyor at the same time. When this happens, the OUTFLOW built-in reports the sum of all inputs exiting the conveyor at that time.

A conveyor with *capacity* specified has a finite capacity and is said to be capacity-constrained, meaning that the sum of its contents cannot exceed the specified capacity. However, specifying capacity in the CONVEYOR function alone will not enforce the constraint, as conveyors will accept whatever input you provide even when a capacity is specified. To make the capacity constraint functional, you must explicitly limit the input values to prevent the conveyor's capacity from being exceeded. Use the MAXINFLOW built-in function to determine the maximum input the conveyor can accept, as follows:

```
input = ...  
limited_input = MIN(input, MAXINFLOW(name))  
name = CONVEYOR(limited_input, tt, cap)
```

Here, the input is prevented from exceeding the value returned by the MAXINFLOW function for this conveyor. This ensures that the conveyor's capacity will not be exceeded.

For a simple example illustrating a conveyor in action, open the "How Do I Use Conveyors" model via **How Do I** under the **Help** menu item.

Ovens


Ovens work a lot like real ovens. They accumulate inputs during the filling state, hold them during the cooking state, and release them during the emptying state. The amount of time the oven spends in the filling and cooking states is controlled by the fill time and cook time parameters.

Ovens are defined using the following syntax:

```
name = OVEN(input, ct, ft, cap)
```

where *name* is the name of the oven, *input* is the input expression, *ct* is the cook time, *ft* is the fill time, and *cap* is the capacity. The name can be followed by array dimensions in square brackets to create an array of ovens. As with conveyors, the oven's capacity argument is optional; if omitted, the oven has unlimited capacity. In a Flowchart, ovens can be used in formula, but they cannot define a reservoir.

The name of the oven is a variable which can be used in other equations or plotted. Its value is sum of all the inputs in the oven multiplied by DT. An example with graphical output is shown in the figure below.



An oven exists in one of four states at any given time: idle, filling, cooking, or emptying. Initially, an oven is in the idle state. While idle, the oven watches its input for nonzero values. The oven will remain in this state indefinitely as long as its input remains zero. When the oven detects a change in the input, it adds that value to its contents and switches to the filling state.

During the filling state, the oven continues adding inputs to its contents. The oven will stop filling once its fill time (sampled when the oven started filling) has expired or its capacity has been reached, whichever comes first. Once this happens, the oven enters the cooking state.

During the cooking state, the oven holds onto its current contents without accepting any additional inputs. It remains in this state until its cook time (sampled when the oven entered the cooking state) has expired. The oven then enters the emptying state.

The emptying state lasts for just one time step. During this step, it empties its contents. This can be observed using the OUTFLOW built-in function whose value will be the sum of all inputs placed in the oven during the filling state:

```
output = OUTFLOW(name)
```

An oven begins filling again during the emptying state if its input expression is nonzero. When this happens, the oven switches back to the filling state for the next time step. Otherwise, it returns to the idle state.

An oven's state can be determined using the OSTATE built-in function:

```
state = OSTATE(name)
```

State values are 0 for idle, 1 for filling, 2 for cooking, and 3 for emptying.

Since an oven cannot always accept inputs, you should use the MAXINFLOW built-in function to limit its input expression. For example, say you have a reservoir of material named "R" which feeds into an oven "V" at the rate of 3 units per time step. You might set up the equations like this:

```
rate = 3  
INIT R = ...  
d/dt(R) = -rate  
V = OVEN(rate, ct, ft, cap)
```

The problem here is that the reservoir will lose material during every time step, even when the oven is cooking. To prevent this, use the MAXINFLOW built-in function to define the amount of material flowing from the reservoir to the oven:

```
rate = 3  
inflow = min(rate, MAXINFLOW(V))  
INIT R = ...  
d/dt(R) = -inflow  
V = OVEN(inflow, ct, ft, cap)
```

During the cooking phase, the MAXINFLOW function evaluates to zero which stops the flow of material out of the reservoir. And while filling, it will prevent the oven's capacity from being exceeded.

For a simple example illustrating oven behavior, open the "How Do I Use Ovens" model via the Help menu in the software.

Queues

Queues accept one or more nonzero inputs during each time step and release them in the order they were added. Unlike conveyors and ovens, the removal of items from a queue is controlled by the use of special built-in functions. This makes it possible for a queue to have multiple outputs. For example, a single queue can feed several conveyors.

Queues are defined using the following syntax:

```
name = QUEUE(input1, input2, ..., inputN)
```

where *name* is the name of the queue and *input1* through *inputN* are the inputs to the queue. The name can be followed by array dimensions in square brackets to create an array of queues.

During each time step, a queue's inputs are evaluated in left-to-right order. If an input is nonzero, it is placed in the queue. Thus, up to N items can be added to a queue during a single time step, where N is the number of inputs.

The name of the queue is a variable which can be used in other equations or plotted. Its value is the sum of all items in the queue multiplied by DT.

By itself, a queue collects inputs but never releases them. To remove items from a queue, use the QPOP and QPOPS built-in functions:

```
output = QPOP(name, maxel, maxamt)
output = QPOPS(name, maxel, maxamt)
```

where *name* is the queue from which elements are removed, *maxel* is the maximum number of items to remove, and *maxamt* is the maximum amount of the items removed. Each time a QPOP or QPOPS built-in is evaluated, it removes as many items as it can from the named queue such that the number of items removed is less than or equal to the *maxel* argument and the sum of the items removed is less than or equal to the *maxamt* argument. Items are always removed in the order they were added.

The QPOP built-in removes only whole items from the queue. If the next item to be removed is greater than *maxamt* argument, no items will be removed and the function's value will be zero.

The QPOPS built-in can remove part of an item from the queue. So, if the next item to be removed is greater than the *maxamt* argument, it will subtract this amount from the item, leaving the remainder in the queue.

Items can be removed from a queue in the same time step that they were added. In other words, material can flow through the queue without spending any time in it. This is in sharp contrast with conveyors and ovens in which inputs are delayed by at least one time step before they exit.

When queues are used to feed conveyors and ovens, the MAXINFLOW built-in function should be used to define the maximum amount of material that can be removed from the queue. For example, a queue feeding an oven would be defined like this:

```
Q = QUEUE(input1, ...)
inflow = QPOP(Q, 1, MAXINFLOW(V))
V = OVEN(inflow, ct, ft, cap)
```

This guarantees that items will be removed from the queue only when the oven is able to accept them. For a more interesting example of the use of queues, open the "How Do I Use Queues" model via the Help menu in the software.

Root Finder Equations

Berkeley Madonna can find roots of one or more expressions automatically during model execution. This feature is useful if certain quantities are defined by zeroing a dependent

expression. For example, say you have a quantity q in your model whose value is determined by solving the following equation:⁵

$$aq^2 + bq + c = 0$$

To determine the quantity of q such that this equation is satisfied, define q using a root finder equation:

```
ROOTI q = a*q*q + b*q + c
GUESS q = 1.0
LIMIT q >= -10
LIMIT q <= +10
```

The ROOTI equation tells Berkeley Madonna to compute the value of q so that the expression $aq^2 + bq + c$ evaluates to zero. Berkeley Madonna uses the GUESS equation to determine the initial value of q when it starts its search for a root. The LIMIT equations specify the range of values q may take during the search. These initial guess and limit equations are required any time you define a root finder equation.

When using a Flowchart model, the root finder equations can be entered in the Globals Window.

Initial and Step Root Finders

The use of the ROOTI keyword indicates that this is an “initial root finder”. Initial root finders perform the root-finding operation only once during model initialization. In the above example, even if a , b , or c were to change during subsequent time steps, the value of q would not be recomputed; it would remain constant through the model’s execution.

If you want the root to be recomputed during each time step, use the ROOTS keywords instead of ROOTI. This defines a “step root finder” which recalculates the root for each time step.

You might think that finding a root during each step of your model would take a long time. However, after the root has been found during initialization, subsequent roots are generally found quickly since the quantities affecting the root (a , b , and c in this example) generally don’t change much. Berkeley Madonna uses the previous root as the initial guess when finding the root the next time and usually doesn’t have to look too far to find it.

The Root Finder Algorithm

Berkeley Madonna finds roots using a globally-convergent variation of the classic Newton-Raphson method. Put simply, it starts with your initial guess and uses Newton’s method to estimate the root’s position. If the function is smooth and the root isn’t too far away, this algorithm converges very quickly. If it doesn’t converge, Berkeley Madonna discards the guess and randomly chooses a new initial guess within the specified limits. While this is somewhat simplistic, it will usually hit upon an initial guess that converges to a root.

⁵ This is a contrived example: you could solve for q using the standard formula for roots of quadratic equations.

How does Berkeley Madonna decide that it has found a root? If the expression on the right-hand side of the root finder equation evaluates to exactly zero, then by definition a root has been found. However, in practice this rarely happens due to the limited precision of floating-point arithmetic on computers. Therefore, Berkeley Madonna also considers a root to have been found when the relative change applied to the variable (as determined by Newton's method) is less than the tolerance specified by the ROOTTOL parameter. In other words, Berkeley Madonna stops when it appears to be close to a root. If it isn't getting close enough for you, try reducing the value of ROOTTOL.

If Berkeley Madonna cannot find a root after a certain number of randomly-chosen initial guesses (currently 1000), it gives up and sets the variable to NAN. This causes a floating-point exception to be reported during model execution. In this example, if no root was found, q would be set to NAN which you could observe by examining its value in a table.

Multidimensional Root Finders

One advantage of using the Newton-Raphson algorithm to find roots is that it scales easily to multiple dimensions. For example, you can set up a series of root finder equations so that the values of x, y, and z are calculated by finding the roots of three separate equations. For example:⁶

$$\begin{aligned}\text{ROOTI } x &= a*x*y + b*x + c \\ \text{ROOTI } y &= d*x*z + e*y + f \\ \text{ROOTI } z &= g*y*y + h*z + i\end{aligned}$$

By examining the interdependencies of these equations, Berkeley Madonna determines that the values of x, y, and z must be determined by simultaneously solving all three equations.

Having three root finder equations in your model does not necessarily mean that you have a three-dimensional root-finding problem. For example:

$$\begin{aligned}\text{ROOTI } x &= a*x*y + b*x + c \\ \text{ROOTI } y &= d*x*x + e*y + f \\ \text{ROOTI } z &= g*y*y + h*z + i\end{aligned}$$

Here, notice that the roots of the first two equations can be found by varying only x and y. Once these roots have been found (a two-dimensional problem), Berkeley Madonna computes the value of z using the third equation without having to vary x or y.

Limit Equations

Limit equations are used to define upper and/or lower bounds for any variable in your model. Whenever a variable with limit(s) is assigned, the value is forced to be within the specified limits. The syntax is:

$$\begin{aligned}\text{LIMIT var } &\geq \text{lower-bound} \\ \text{LIMIT var } &\leq \text{upper-bound}\end{aligned}$$

⁶ The initial GUESS and LIMIT equations have been omitted for clarity, but they must be included.

“var” must be the name of a variable defined elsewhere in your model. “lower-bound” and “upper-bound” may be any valid Berkeley Madonna expression. On macOS, you can use \geq and \leq instead of $>=$ and $<=$, respectively.

As with ROOTI/ROOTS, the LIMIT function can be added to the Globals Window when using the Flowchart.

Other Equation Types

There are three more types of “equations”. Unlike other equations, however, they don’t directly affect the values of variables in your model, so we call them “statements” instead.

METHOD Statement

The “method” statement defines the default integration method that appears in the parameter window when your model is first compiled.⁷ The syntax is:

METHOD name

“name” must be Euler, RK2, RK4, Auto, or Stiff.

DISPLAY Statement

The “display” statement allows you to control which symbols in your model are accessible from various windows and dialogs. The syntax is:

DISPLAY name1, name2, ...

“name1”, “name2”, ... are names of symbols in your model. You can use more than one display statement if you want. If your model uses one or more display statements, any symbols whose names don’t appear in one of these statements will not be available in any windows or dialogs. If you don’t use any display statements, all your symbols will be available.

For visual models, add the DISPLAY command in the Globals Window.

RENAME Statement

The “rename” statement enables you to rename the built-in symbols listed in the next section. To rename a built-in symbol, use the following syntax:

RENAME old-name = new-name

For example, to change the name of TIME to X, you would write:

RENAME TIME = X

Note that the new name must not already be in use. For example, you cannot rename TIME to STOPTIME (unless STOPTIME itself has already been renamed to something else).

⁷ Berkeley Madonna remembers the current integration method whenever you save your model. So, the METHOD statement only has an effect for models that have never been compiled and saved.

For visual models, use the RENAME command from the Globals Window.

Built-in Symbols

The following symbols are implicitly defined in all models. You can refer to them in your equations just like any other variables in your model. Their values can be explicitly specified in your equations if the default shown below is not appropriate. Note that the definition of TIME cannot be changed.

Symbol	Default Definition
STARTTIME	STARTTIME = 0
STOPTIME	STOPTIME = 10
DT	DT = 0.02
DTMIN	DTMIN = 1e-6
DTMAX	DTMAX = 1.0
TOLERANCE	TOLERANCE = 0.01
DTOUT	DTOUT = 0
ROOTTOL	ROOTTOL = 0.001
TIME	INIT TIME = STARTTIME d/dt(TIME) = 1.0

DT is used only by fixed-stepsize integration methods. DTMIN, DTMAX, and TOLERANCE are only used by variable-stepsize integration methods. Symbols that are not applicable to the currently-selected integration method are not shown in the parameter window. See Integration Methods on page 45 for more information.

DTOUT controls the interval at which results are stored in memory. See Using DTOUT on page 47.

ROOTTOL appears only when your model uses one or more built-in root finder equations. See Root Finder Equations on page 30.

Operators

The following table lists the arithmetic, logical, and relational operators supported by Berkeley Madonna. They are listed in order of decreasing precedence. The relational operators ($=$, \neq , etc.) generate a value of 1 if the relation is true or 0 if the relation is false. The logical operators (AND, OR, NOT) interpret zero as false and all other values, even those very close to zero, as

A warning on operator precedence

The binding of unary minus is non-standard and will produce incorrect results if used in the normal mathematical sense. Due to historical reasons, in Berkeley Madonna, the expression $-a^2$ is equal to $(-a)^2$ since $-$ binds tighter than $^$. Use $-(a^2)$ instead. For example, in Berkeley Madonna: $-3^2 = 9$ and $-(3^2) = -9$.

true. The conditional operator (IF-THEN-ELSE) returns the value of the “then” expression if the “if” expression is nonzero; otherwise, it returns the value of the “else” expression.

Operator	Name	Example
+ - NOT	unary plus (no effect) negate logical inverse	+a -b NOT c
^, **	raise to power	a ^ b, a ** b
* /	multiply divide	a * b a / b
+ -	add subtract	a + b a - b
= ≠, <> < > ≤, <= ≥, >=	equal to not equal to less than greater than less than or equal to greater than or equal to	a = b a ≠ b, a <> b a < b a > b a ≤ b, a <= b a ≥ b, a >= b
AND	logical and	a AND b
OR	logical or	a OR b
IF-THEN-ELSE	conditional	IF a THEN b ELSE c
INF, ∞	infinity	z = a / INF

Built-in Functions

Berkeley Madonna supports the following built-in functions:

Function	Meaning
ABS(x)	Absolute value of x
ARCCOS(x)	Inverse cosine of x [domain: -1 ≤ x ≤ 1]
ARCCOSH(x)	Inverse hyperbolic cosine of x [domain: x ≥ 1]
ARCSIN(x)	Inverse sine of x [domain: -1 ≤ x ≤ 1]
ARCSINH(x)	Inverse hyperbolic sine of x
ARCTAN(x)	Inverse tangent of x
ARCTAN2(y,x)	Angle (radians) in the Euclidian plane between positive

Function	Meaning
	x-axis and ray going to point (x,y)
ARCTANH(x)	Inverse hyperbolic tangent of x [domain: $-1 < x < 1$]
ARRAYSUM(x[*])	Sum of elements in array x
ARRAYMEAN(x[*])	Average of elements in array x
ARRAYSTDDEV(x[*])	Standard deviation of elements in array x
BINOMIAL(p, n) BINOMIAL(p, n, s)	Binomial random number generated from n trials of probability p, optional seed s
COS(x)	Cosine of x
COSH(x)	Hyperbolic cosine of x
DELAY(x, d) DELAY(x, d, i)	Value of x delayed by time d, optional initial value i
ERF(x)	Error function of x
ERFC(x)	Complimentary error function of x
EXP(x)	e (2.71828...) raised to x
GAMMALN(x)	Natural log of the gamma function of x [domain: $x \geq 0$]
GRAPH(v) (x1,y1) (x2, y2) ...	Graphical function of v, with one or more (x,y) data points.
INIT(x)	Value of expression x at initialization.
INT(x)	Integer part of x
LOG10(x)	Log base 10 of x [domain: $x > 0$]
LOGN(x)	Natural log of x [domain: $x > 0$]
MAX(x1, x2, ...)	Maximum of x1, x2, ...
MAXINFLOW(s)	Maximum allowable inflow into conveyor or oven s
MEAN(x1, x2, ...)	Average of x1, x2, ...
MIN(x1, x2, ...)	Minimum of x1, x2, ...
MOD(x, y)	Floating-point remainder of x / y. This is equal to: $x - y * \text{INT}(x / y)$
NETFLOW(r)	Flow into reservoir r from previous time step
NORMAL(μ , v) NORMAL(μ , v, s)	Normally-distributed random number with mean μ and variance v, optional seed s

Function	Meaning
OSTATE(s)	State of oven s: 0=idle, 1=filling, 2=cooking, 3=emptying
OUTFLOW(s)	Outflow from conveyor or oven s
PI	Value of π (3.14159...)
POISSON(μ) POISSON(μ , s)	Random integer drawn from a Poisson distribution with mean μ , optional seed s
PULSE(v, f, r)	Pulse with volume v, first pulse at time f, repeat interval r
QPOP(q, n, m)	Removes up to n items from queue q totaling m or less
QPOPS(q, n, m)	Like QPOP, but can remove part of an item to satisfy m
RANDOM(l, h) RANDOM(l, h, s)	Uniformly-distributed random number between l and h, optional seed s
ROUND(x)	x rounded to the nearest integer
SIN(x)	Sine of x
SINH(x)	Hyperbolic sine of x
SQRT(x)	Square root of x [domain: $x \geq 0$]
SQUAREPULSE(t, d)	Square pulse of height 1 starting at time t and lasting for duration d
STEP(h, t)	Step of height h at time t
STEPSIZE	Change in TIME from previous step to current step
SUM(x1, x2, ...)	Sum of x1, x2, ...
SWITCH(x, y)	Equivalent to IF $x > y$ THEN 1 ELSE 0
TAN(x)	Tangent of x
TANH(x)	Hyperbolic tangent of x

Some built-in functions are defined on a restricted domain, such as the inverse trigonometric functions, and when called outside this domain it will raise an exception. How Berkeley Madonna handles this depends on your preference settings. If **Stop on Exception** under **Model/Settings...** is selected, the run will terminate with an error; however, if **Stop on Exception** is deselected, then the model will continue through errors, which depending on the model may or may not produce useful data after the exception. The domain is provided in the table above for Built-in functions with restrictions.

Most of these functions are self-explanatory. A few, however, deserve some explanation.

ARRAYSUM, ARRAYMEAN, ARRAYSTDDEV

These functions require a single array argument. To pass an array “v” to one of these functions, use the syntax “v[*]”. For example:

```
v_sum = ARRAYSUM(v[*])  
v_avg = ARRAYMEAN(v[*])  
v_std = ARRAYSTDDEV(v[*])
```

For more information on arrays, see Array Equations on page 40.

BINOMIAL, NORMAL, POISSON, RANDOM

These functions generate various distributions of random numbers using a lagged Fibonacci generator with a very long period (approximately 10^{84}). This generator is based on the [FIBMULT](#) program. Each function comes in short and long forms. The short form uses a different seed each time you run your model. The long form uses a fixed seed value that you supply. When using a fixed seed, you’ll get a repeatable sequence of random numbers each time you run your model (as long as you don’t change the seed).

DELAY

This function returns its input value delayed by the specified delay time. The delay time expression (2nd argument) is evaluated once during model initialization so the amount of delay remains fixed for the entire run. The function returns a special “initialization” value until TIME has reached the specified delay time. In the two-argument form of the function, the initialization value is the initial value of the input expression. In the three-argument form, the initialization value is specified by the third expression which, of course, is evaluated only once.

The DELAY function works properly for both fixed- and variable-stepsize integration methods. However, if an expression using DELAY is used as a reservoir’s flow equation, the net flow contributed by DELAY will not be affected by the integration method. In other words, when integrating expressions using DELAY, the accuracy of those reservoirs will be that of Euler’s method no matter what integration method is actually used. The reason is that the internal queue that DELAY uses is not updated during the trial steps used by these higher-order methods.

NETFLOW

This function takes a single reservoir as an argument and returns the amount the reservoir changed from the previous time step to the current step. It returns zero during the initialization step. It is especially useful in conjunction with the Custom DT Method (described on page 48) for controlling the stepsize in terms of reservoir changes. Note that you can’t simply multiply a reservoir’s flow expression by DT to compute the net change unless you’re using Euler’s method.

PULSE

PULSE instantaneously changes time-dependent variables by adding a fixed amount to a quantity over a short amount of time. These pulses are generally sharp when using fixed time step methods (Euler, RK2, RK4); however, when used in combination with variable time step methods, the pulse is often smoothed out over the transition in a manner that may not be consistent with the user’s desired outcome. To sharpen the pulse, reduce DTMAX and visually

examine the shape of the pulsed quantity until it produces the desired result. Alternatively, users may want to default to fixed timestep method when employing these functions.

Berkeley Madonna employs a piece wise linear hat function to calculate the pulse function. Some common things that you should know about using Pulse are:

- 1) The volume v is added over 2-3 timesteps.
- 2) For variable time step methods, limit DTMAX to a small non-zero number to control how quickly the volume is added. For fixed timestep methods, use DT to control how quickly the volume is added.
- 3) To ensure that the full pulse v is added every repeat interval r for $\text{pulse}(v, f, r)$, make sure that the repeat time r is several times longer than DT for fixed time step methods. Similarly make sure the repeat time r is several times longer than DTMAX for variable time step methods.

SQUAREPULSE

The SQUAREPULSE(t,d) function uses conditional statements to instantaneously change the value of a function at time t lasting for duration d ; however, if the timing of the change occurs within a large time step (for long DT values compared to d), then the quantity may not change. Hence, as with pulse, care must be given whether using variable or fixed time step solvers to ensure the square pulse behaves as expected.

STEPSIZE

This function returns the time difference between the previous step and current step (zero during initialization). This is useful for monitoring the change in stepsize when using variable-stepsize methods: assign the value of STEPSIZE to a variable and then add this variable to a graph window.

MAXINFLOW, OUTFLOW, OSTATE, QPOP, QPOPS

These functions are used in conjunction with conveyors, ovens, and queues. See Discrete Equations on page 27 for details.

GRAPH

This function allows the user to define a piece-wise linear function by entering a set of paired points. These values can be entered directly into the Equations Window or into the Globals Window for visual models using syntax similar to the following examples. The following usage creates a peaked function as a function of time that goes from 0 to 8 and then back to zero reaching a maximum at TIME = 3:

$$y = \text{graph}(\text{TIME}) (0,0) (1,1) (2,3) (3,8) (4,3) (5,1) (6,0)$$


For values of time greater than 6, y will take on the last value of zero. In a second example, y is a saw toothed function of TIME that starts at -3, reaches +3 at TIME = 1, and then back to -3 at TIME = 2:

$$y = \text{graph}(\text{mod}(\text{TIME},2)) (0,-3) (1,3) (2,-3)$$

The mod function makes y repeat this pattern again with a period of 2. In either case, y is now a function that can be used in your models just as any other functions can be used.

The GRAPH function cannot be used in the Flowchart, but the **Create Graph** button found within the **Flow** and **Formula** (not **Reservoirs**) dialog windows provides the same functionality. First, open the dialog for a formula or flow, in this case formula F1 (see figure):

- 1) Set the equation equal to the graph's independent variable: $F1 = \text{TIME}$, if you want a function of time, or $F1 = R1$, if you want the graph to change as a function of reservoir R1 (there must be an arc from R1 to F1 in this case).
- 2) Select **Create Graph** (or **Edit Graph** if you previously created the Graph). Enter the x and y minimum/maximum values, the number of data points N (must be 6 or greater), and then manually enter the y values into the far right column (see left panel in figure). Clicking and dragging the graph (blue line) to set the values does not currently work. Alternatively, you can load a file containing the data points with the **Load Data** button.
- 3) Click OK and then exit the formula window and look at the equations window. You should see $F1 = \text{GRAPH}(\text{TIME})$, or $F1 = \text{GRAPH}(R1)$, and running the model should produce a graph like the right panel in the figure for the data given on the left.
- 4) Now use the formula F1 in the flowchart as usual.



Array Equations

Berkeley Madonna supports array equations in addition to scalar equations. This feature can only be used in the Equations Window when there is no graphical Flowchart or in the Globals Window of a graphical model – there is no support for arrays in the Flowchart Window itself. Arrays are useful if you have many scalar equations with similar right-hand sides. Arrays can have one, two, or three dimensions. To define an array equation, the variable name on the left-hand side includes a range of subscripts enclosed in square brackets like this:

$$v[1..10] = \text{SIN}(2*\text{pi}*f)$$

In this example, v is defined as a vector (one-dimensional array) of 10 elements with subscripts 1 through 10. Note that you don't need to explicitly define the number of elements in the array; Berkeley Madonna figures out the size based on the subscript range on the left-hand side.

Arrays with two or three dimensions are defined with two or three subscript ranges. For example, the following defines a two-dimensional array of 60 elements (5 rows by 12 columns):

$$a[1..5,1..12] = 10.5$$

And the following defines a three-dimensional array with 24 elements:

$$a[1..2, 0..2, 2..5] = -77$$

This last example is a bit tricky because index for each dimension starts at a different number. The array's dimensions are 2 x 3 x 4 and the elements are a[1,0,2], a[1,0,3], a[1,0,4], a[1,0,5], a[1,1,2], ..., a[1,1,5], ..., a[2,0,2], ..., a[2,2,5].

You can define a single array using more than one equation:

$$\begin{aligned} k[0..3] &= a * \text{COS}(x) \\ k[4..9] &= b * \text{SIN}(x) \\ k[10] &= 0.5 \end{aligned}$$

In this example, k is an 11-element vector with subscripts 0 through 10. When multiple equations are used to define an array, they are executed in the order they appear in the equation file. Within a single equation, the elements are assigned from the first subscript to the last. In the previous example, the order of assignment is k[0], k[1], k[2], ..., k[9], k[10]. This execution order can very important as we shall soon see.

Berkeley Madonna defines special indexing variables named **i**, **j**, and **k** which are valid only in the context of an array equation. They refer to the subscript of the element being assigned on the left-hand side. Most often, these variables are used to refer to array elements. For example:

$$\begin{aligned} a[1..10] &= ... \\ b[1..10] &= ... \\ \text{total}[1..10] &= a[i] + b[i] \end{aligned}$$

In two- or three-dimensional array equations, there are two or three index variables. Each index variable corresponds to one dimension: **i** to the first dimension, **j** to the second, and **k** to the third. For example, the in following three-dimensional array equation, **i** goes from 1 to 3, **j** goes from 1 to 4, and **k** remains unchanged at 5:

$$s[1..3,1..4,5] = f[i,j] * g[j,k]$$

The rightmost index variable (**j** in a two-dimensional equation, **k** in a three-dimensional equation) changes the fastest. For example, in the following equation:

$$a[1..3,1..3] = ...$$

the elements are assigned in the following order:

$$\begin{aligned} a[1,1] &= ... \\ a[1,2] &= ... \\ a[1,3] &= ... \\ a[2,1] &= ... \\ a[2,2] &= ... \\ a[2,3] &= ... \\ a[3,1] &= ... \end{aligned}$$

a[3,2] = ...
a[3,3] = ...

Berkeley Madonna allows you to refer to elements of the array being assigned in the right-hand side of its equation. This permits you to write equations like this:

factorial[1] = 1
factorial[2..30] = factorial[i - 1] * i

In this example, the order of evaluation is extremely important. If the two equations were reversed, factorial[2] would be undefined because the value of factorial[1] has not yet been computed. Another situation to avoid is illustrated here:

a[10] = 1000
a[1..9] = a[i+1] * 0.97

This example is invalid because the computation of a[1] depends on the value of a[2] which hasn't been computed yet. The correct way to write these equations is:

a[10] = 1000
a[9..1] = a[i+1] * 0.97

Finally, you must watch out for gaps in your arrays. Consider the following:

b[0] = 0
b[1..4] = b[i-1] + sqrt(i)
b[10] = 100
b[9..6] = b[i+1] * b[10-i]

Everything looks fine here until you realize that b's subscripts go from 0 to 10, but b[5] is not defined anywhere. Berkeley Madonna does not check for such errors, so use care.

Arrays can be used with differential, difference, and discrete equations. Just remember to include the subscript range immediately following the variable name on the left-hand side. For example:

y[1](starttime) = 1
y[2..3](starttime) = 1.1 * y[i - 1]
d/dt(y[1..3]) = -0.1 * y[i]

INIT a[1..5] = i
NEXT a[1..5] = a[i] + i

INIT x[1..n] = 0
x[1..n](t) = x(t - dt) + (1 / sqrt(y[i])) * dt

c[1..n] = CONVEYOR(...)

Note that you can't use multiple equations to define an array of discrete objects. This restriction is necessary to prevent the creation of schizophrenic arrays such as this:

s[1..3] = CONVEYOR(...)
s[4..6] = QUEUE(...)
s[7..9] = OVEN(...)

Limits can be applied to arrays, but the limit applies equally to all elements. You cannot apply a limit to only certain elements. For example, to limit each element of vector y in the previous example to values greater than or equal to 0.5, add the following equation:

$$\text{LIMIT } y \geq 0.5$$

You can use variables in subscript ranges. This capability makes it easy to change the size of arrays without recompiling your model. For example:

```
factorial[1] = 1
factorial[2..n] = factorial[i-1] * i
n = 10
```

Since n is a parameter, you can change its value from the parameter window. Berkeley Madonna dynamically determines the size of the array before each run and allocates the necessary memory.

If a subscript range uses an expression whose value changes over time, the size of the array does not change. The expression's value during the initialization step determines the size of the array and this size remains fixed during the run.

Dataset Functions

You can use imported datasets as piecewise-linear functions in your model's equations. This feature makes it easy to use data from external sources to control your model's behavior. Before you can use a dataset function in your equations, you need to import a dataset, which must be in ASCII format. Importing Datasets on page 16 describes how to do this.

Once you've imported a dataset, you can use it just like a built-in function taking one argument (for vector datasets) or two arguments (for matrix datasets). For example, assume you've imported a vector dataset named "#temperature" which maps distances into temperature values. (This could be from taking temperature measurements at various points along a metal bar.) In your model, you could compute the temperature " t " at a point " p " using the following expression:

$$t = \#temperature(p)$$

Note that since dataset names always begin with a pound-sign character ($\#$), they can be easily distinguished from regular built-in functions. To evaluate this expression, Berkeley Madonna looks up the value of " p " in the dataset's array of input values (X). It finds the two X values which bracket the input value, then it uses linear interpolation to compute the output value using the corresponding Y values. For values outside the dataset's range of X values, it uses the Y value corresponding to the closest X value (always the first or last point in the dataset).

Matrix datasets are used in a similar fashion except that they take two arguments. For example, assume we have a matrix dataset representing temperatures measured at various points on a metal surface. You can determine the temperature at a point with Cartesian coordinates " x " and " y " like this:

$$t = \#temperature(x, y)$$

As with vector datasets, Berkeley Madonna performs linear interpolation to find the output value for the corresponding inputs. This time, the interpolation is performed in two steps since both the X and Y inputs may not correspond exactly to the input values in the dataset.

Plotting Matrix Datasets

As discussed in Matrix Datasets on page 17, matrix datasets cannot be directly plotted in a graph window as vector datasets can. However, you can still get an idea of what a matrix dataset looks like by using an array equation to sample its outputs for various inputs and then plotting the resulting array. To do this, you need to know the range of the dataset's input domains (the X and Y inputs). Then, construct an array equation that takes output values at various inputs over the domain. One of the input domains is covered using the TIME variable and the other is covered by the array index variable *i*.

For example, assume you have a matrix dataset named "#temperature" that maps points in a plane to temperature values. Assume the X coordinates go from 0.0 to 5.0 and the Y coordinates from 0.0 to 10.0. Sample the elements of this dataset using the following equations:

```
STARTTIME = 0
STOPTIME = 10.0
DT = 1.0
T[0..5] = #temperature(i, TIME)
```

Run the model and plot the value of array T. You'll get six curves representing temperatures in the Y direction for six different X values. Of course, you could instead use the following equations:

```
STARTTIME = 0
STOPTIME = 5.0
DT = 1.0
T[0..10] = #temperature(TIME, i)
```

This would result in eleven curves representing temperatures in the X direction for eleven distinct Y values.

Tips and Techniques

Creating Periodic Functions

To make any function repeat periodically, use the MOD built-in. If F is a function that you want to repeat over the time interval I, you would write:

$$y = F(\text{MOD}(\text{TIME}, I))$$

For example, to make a square wave with a period of one time unit and an amplitude of 1, you could write:

$$y = \text{IF } \text{MOD}(\text{TIME}, 1) \geq 0.5 \text{ THEN } 1 \text{ ELSE } 0$$

or, more succinctly:

$$y = \text{MOD}(\text{TIME}, 1) \geq 0.5$$

Creating Piecewise-Linear Functions

Multiple IF-THEN-ELSE expressions can be used to write piecewise linear differential equations. For example:

$$d/dt(x) = \text{IF } x < 3 \text{ THEN } L1 \text{ ELSE IF } x \leq 4 \text{ THEN } L2 \text{ ELSE } L3$$

where L1, L2, and L3 are line segments. Or you can use the GRAPH() built-in which generates the line segments for you:

$$d/dt(x) = \text{GRAPH}(x) (0,0) (3,4) (4,6.5) (10,1)$$

For large numbers of line segments, you may find it easier to list the points in a text file and use that file to define a dataset function. See **Dataset Functions** on page 43 for details.

Non-negative Reservoirs and Unidirectional Flows

Berkeley Madonna assumes that reservoirs and flows can take any real value.

If your model depends on reservoirs and/or flows being limited to non-negative values, you'll need to add appropriate LIMIT statements to ensure these conditions are met. For example, the following equation taken from STELLA defines a reservoir named "Resv" which depends on an inflow named "In" and an outflow named "Out":

$$d/dt(\text{Resv}) = \text{In} - \text{Out}$$

By default, Berkeley Madonna assumes that "Resv", "In", and "Out" can be positive or negative. To make "Resv" non-negative and the two flows unidirectional, add the following statements to the equation file:

```
LIMIT Resv >= 0
LIMIT In >= 0
LIMIT Out >= 0
```

If your model depends on STELLA's "outflow prioritization" mechanism, you'll have to rewrite your equations in Berkeley Madonna to duplicate this behavior.

Integration Methods

Fixed-stepsize Methods

Berkeley Madonna provides three fixed-stepsize integration methods: Euler's method, Runge-Kutta 2, and Runge-Kutta 4. These methods step the model in increments of DT. The number of steps (not including the initialization step) is equal to: (STOPTIME - STARTTIME)/DT.

Variable-stepsize Methods

Berkeley Madonna also provides two variable-stepsize integration methods. These algorithms choose the largest stepsize they can consistent with the tolerance and minimum/maximum stepsize you specify. The DTMIN parameter specifies the minimum allowable stepsize and the size of the first step taken. The DTMAX parameter specifies the maximum allowable stepsize (it also controls the width of impulses generated by the PULSE built-in). The variable stepsize algorithms automatically adjust the stepsize within these limits to minimize the number of steps taken while keeping the estimated error below the threshold specified by the TOLERANCE parameter.

When using a variable-stepsize method, you may want to monitor how the stepsize is changing as the run progresses. Your first inclination might be to simply plot DT. However, you won't find DT among the list of available variables to plot. That's because DT is a parameter whose value by definition is constant throughout the run. When using variable

stepsize methods, the value of DT is not used and therefore it is not shown in the parameter window. The actual stepsize is determined using the STEPSIZE built-in function. For example:

$$h = \text{STEPSIZE}$$

By assigning the value of the STEPSIZE built-in to a variable, you can monitor it during the run by simply plotting the variable (“h” in this example).

Here is the process used by the variable-stepsize algorithms to advance the model by one time step:

1. Compute flows for each reservoir in the model based on the current values of all variables.
2. Estimate absolute error in each flow computed in step 1.
3. Convert absolute errors computed in step 2 to relative errors by dividing each error by the magnitude of the reservoir’s current value. This scaling is not performed if the magnitude of the reservoir is less than 1.
4. If the maximum relative error is greater than TOLERANCE and the stepsize is greater than DTMIN, reduce the stepsize and go back to step 1.
5. Advance the model to the next time step using the current stepsize.
6. If the maximum relative error is less than TOLERANCE and the stepsize is less than DTMAX, increase the stepsize for the next time step. Otherwise, leave the stepsize unchanged.
7. Return to step 1 to do the next time step.

The computation of the estimated error for each reservoir can be stated formally:

Let y be the value of the reservoir, Δy be the estimated increment (flow) to be added to the reservoir, and ε be the estimate of the error in Δy . The estimated error is:

$$\left\{ \begin{array}{l} \left| \frac{\varepsilon}{y} \right|, |y| > 1 \\ |\varepsilon|, |y| \leq 1 \end{array} \right\}$$

The Auto-stepsize and Rosenbrock (stiff) methods differ in the way they compute flows, estimate absolute errors, and adjust the stepsize.

The Auto-stepsize method uses a fifth-order Runge-Kutta algorithm to compute flows. It estimates errors by comparing this flow with another flow computed using a fourth-order algorithm. The difference between these estimated flows is the estimated error.

The Rosenbrock (stiff) method uses a semi-implicit fourth-order Runge-Kutta algorithm to compute flows. It estimates error by comparing this flow with another flow computed using a third-order algorithm. Again, the difference between these estimated flows is the estimated error.

The Auto-stepsize method adjusts the stepsize by multiplying it by the following factor:

$$0.99 \times \left(\frac{tol}{\epsilon_{\max}} \right)^{1/5},$$

where *tol* is the TOLERANCE and ϵ_{\max} is the maximum relative error.

This factor is limited to the range 0.1 - 10; thus, the stepsize cannot change by more than a factor of ten from one step to the next. The fifth root is used because, to a first approximation, the change in the error is proportional to the change in stepsize raised to the fifth power for a fourth-order method. The 0.99 is a “safety factor”: it’s much better to use a slightly smaller stepsize than theoretically possible; otherwise, the algorithm may end up overestimating the stepsize slightly and consequently discarding attempted steps (wasting time) because stepsize was just a little too large.

The Rosenbrock (stiff) method adjusts the stepsize in the same way except that a different factor is used:

$$0.99 \times \left(\frac{tol}{\epsilon_{\max}} \right)^{1/4}$$

This factor is limited to the range 0.5 - 2; thus, the stepsize cannot change by more than a factor of two from one step to the next. The fourth root is used because the change in error is roughly proportional to the change in stepsize raised to the fourth power for a third-order method.

For stiff systems of equations, the Auto-stepsize method grossly overestimates errors and thus uses much smaller stepsizes than is actually necessary. The Rosenbrock (stiff) method does a much better job for these types of systems.

On the other hand, for smooth, non-stiff systems, the Auto-stepsize method, being of higher order than the Rosenbrock (stiff) method, can take larger steps for a given TOLERANCE. Also, the Rosenbrock (stiff) method does a lot more math (matrix inversion, etc.) to estimate flows. For these systems, the Auto-stepsize algorithm is a better choice.

The Auto-stepsize algorithm is derived from the routine **rkqs()** published in *Numerical Recipes in C*. The Rosenbrock (stiff) method is derived from the routine **stiff()** in the same text.

Using DTOUT

Normally, Berkeley Madonna stores an output point for each time step. For small stepsizes, this can result in large quantities of output data. You can reduce the amount of data by increasing the stepsize, but often this is not possible because a small stepsize is needed to achieve sufficient accuracy. By using DTOUT, you can control how much output data are stored independently of the stepsize used to compute them.

The default value of DTOUT is zero, which causes Berkeley Madonna to store every computed step in memory. To reduce the number of steps stored, set DTOUT to some multiple of the stepsize. For example, if your model uses a fixed-stepsize integration method and a DT of 0.001 and you want to store only one output point for every 100 computed points, set DTOUT to 0.1 (0.001 * 100). The resulting output will appear to have a DT of 0.1, but the data are still computed with a DT of 0.001.

If you are using a fixed stepsize integration method and DTOUT is not an integer multiple of DT, Berkeley Madonna will round DTOUT to the nearest integer multiple of DT and use this

rounded value in lieu of the value you specify. For example, if your DT is 0.004 and you set DTOUT to 0.05, Berkeley Madonna will round DTOUT to 0.048 and store 1 output step for every 12 computed steps.

Variable stepsize methods handle DTOUT differently. When each computed step is finished, Berkeley Madonna checks to see if the current time is at least DTOUT greater than the time of the previous output step. If it is, the current step will be stored. Otherwise, it is skipped. Note that Berkeley Madonna will always store the final step when TIME reaches STOPTIME, so the interval between the last two output steps may be less than DTOUT (or DTMIN for that matter).

It's important to remember that DTOUT affects only the number of steps stored in memory. Berkeley Madonna still uses the specified stepsize to perform all internal calculations, so the accuracy of the solution is not affected as evidenced by the fact that the number of computed steps shown in the upper right-hand corner of the graph window doesn't change.

Custom DT Method

Berkeley Madonna provides a custom DT method which allows your model to directly control the value of DT used for each step. To use this feature, define DT as a difference equation in the following format:

```
INIT DT = initial_dt
NEXT DT = IF ok_to_step THEN next_dt_value ELSE retry_dt_value,
```

where initial_dt, next_dt_value, retry_dt_value are all numeric quantities (constants or functions), and ok_to_step is 1 or 0 based on a conditional statement.

When you define DT this way, Berkeley Madonna runs your model using the following steps:

1. Initial values of all reservoirs and difference equations are computed as usual. The value of DT is set to initial_dt.
2. Flows are computed for each reservoir based on the current value of DT.
3. The ok_to_step expression is evaluated. If this expression is true (nonzero), Berkeley Madonna continues with step 4. Otherwise, DT is changed to retry_dt_value and Berkeley Madonna returns to step 2.
4. The model is advanced to the next time step using the current value of DT.
5. DT is changed to next_dt_value and Berkeley Madonna goes to step 2.

Note that DT is not limited to any minimum or maximum value; it is your responsibility to set DT to appropriate value in the next_dt_value and retry_dt_value expressions. DTMAX is used only to determine the width of the PULSE function.

Here is an example problem: Vm is a reservoir representing voltage. You want to step your model as fast as possible with the limitation that the change in the voltage cannot exceed ± 10 millivolts per time step. Also, DT must be limited to 8 milliseconds maximum and 1 microsecond minimum. When the proposed change in Vm is too large, you'll try reducing DT by a factor of 2; when the change is within bounds, try increasing DT by 10%.

Here's what you need to put in your model to accomplish the above:


```
INIT DT = .001      {Start with 1 millisecond time steps}
NEXT DT = IF ABS(NETFLOW(Vm)) <= .010 THEN MIN(DT * 1.1, 0.008) ELSE MAX(DT
/ 2, 1e-6)
```

Note that if the minimum DT of 1e-6 produces too large a voltage change, the model will loop infinitely. This should be avoided by changing the ok_to_step expression to:

```
ABS(NETFLOW(Vm)) <= 0.010 OR DT = 1e-6
```

This way, the model is guaranteed to step when trying the minimum allowable DT.

One word of caution: you may be tempted to use intermediate variables to reduce the complexity of your ok_to_step and retry_dt_value expressions. For example:

```
ok_to_step = ABS(NETFLOW(Vm)) <= 0.010 OR DT = 1e-6
NEXT DT = IF ok_to_step THEN ... ELSE ...
```

This will not work! Why? Because the ok_to_step expression is evaluated before your model is advanced. So, the ok_to_step variable above will still have its value from the previous step which depends on the flow from the previous step. To avoid problems like this, you should make sure that your ok_to_step and retry_dt_value expressions do not reference variables that depend on DT or the NETFLOW() function. Instead, references to DT and NETFLOW() should appear directly in these expressions.

Another thing to note: when Berkeley Madonna retries a step because the ok_to_step expression was false (zero), the only thing that changes is DT. Nothing else in your model has changed because no step was taken. Thus, the only things that will be different when the ok_to_step expression is re-evaluated are DT and the flows going into your reservoirs (which depend on DT and are computed using the NETFLOW() function).

Models using the custom DT feature can employ any fixed-stepsize integration method (Euler's method, Runge-Kutta 2, and Runge-Kutta 4). Variable-stepsize algorithms are not supported. When using the custom DT feature, your model is implicitly running with a variable stepsize.

Other Features

Chemical Reactions

The chemical reactions feature enables you to enter equations using ordinary chemical reaction notation. Berkeley Madonna generates the corresponding kinetic equations based on the law of mass action for you. Chemical reactions are represented by special embedded objects in the equation window which makes them easy to identify. You can freely mix these objects with your own equations.

To work with chemical reactions in your model, choose **Chemical Reactions** from the **Modules** submenu in the **Model** menu. Berkeley Madonna displays the **Chemical Reactions** dialog. You use this dialog to add, modify, and delete chemical reactions in your model.

To add a chemical reaction to your model, enter the left- and right-hand sides of a reaction into the **Reactants** and **Products** fields, respectively. Each side consists of a list of substances (reactants or products) separated by plus signs (+). Each substance is prefixed with an optional integer multiplier. For example, the following reaction consumes one unit of substance A and two units of substance B, producing three units of substance C and one unit of substance D:



Enter the forward and reverse rate constants into the **Kf** and **Kr** fields, respectively. Then click the **Add** button. The chemical reaction is added to the **Reactions** list. Note that each reaction is assigned a number so it can be uniquely identified. Also note that the forward and reverse reaction rates are not shown in the list. To see them, select the reaction and examine the **Kf** and **Kr** fields. You can continue adding reactions to the list by repeating the above process.

To modify an existing reaction, select it in the **Reactions** list, edit the **Reactants**, **Products**, **Kf**, and/or **Kr** fields and then click the **Modify** button. Don't forget to click **Modify** or the reaction will not be updated. To remove a reaction, simply select the reaction and click **Remove** (or double-click the reaction).

As you add reactions, you'll notice that Berkeley Madonna adds entries to the **Initial Concentrations** list. Here you can specify the initial value of any substance. To do so, select the substance and enter the value in the box below the list.

When you're done building the list of reactions and setting the initial concentrations for each substance, click the OK button. Berkeley Madonna adds the chemical reactions to your model's equations. To see the results, open the equation window by choosing **Equations** from the **Model** menu.

Each chemical reaction is represented in the equation window by an embedded chemical reaction object. These objects are easily identified by their gray shading. Within each object, Berkeley Madonna shows the chemical reaction in chemical notation followed by the reaction rate equation (RXN1 = ...), forward and reverse rate constants, initial concentrations, and the differential equations for each substance.

You can simplify the display of chemical reaction objects by hiding any of these parts except for the reaction itself by unchecking the appropriate **Show** boxes in the **Chemical Reactions** dialog. Note that these settings apply to all reactions in your model.


When you compile a model with chemical reactions, Berkeley Madonna generates kinetic equations for each reaction as shown in its object. Berkeley Madonna keeps track of which substances are produced and/or consumed by each reaction, keeping the generated equations consistent with each other. To see the actual equations generated, activate the equation window (Model/Equations).

To quickly modify a reaction, double-click it in the equation window. Berkeley Madonna will open the **Chemical Reactions** dialog and select the desired reaction for you. All you have to do is enter the changes, click **Modify**, then click OK. Note that the default buttons are set up so that you can just hit enter twice rather than clicking the buttons.

You don't have to keep the equation window open to modify chemical reactions. You can always choose **Chemical Reactions** from the **Modules** submenu in the **Model** menu to display the dialog.

Adding Other Equations

As mentioned earlier, you can freely mix your own equations with chemical reaction objects. This can cause problems if your equations conflict with those generated by the chemical reaction objects. For example, consider the following simple reaction:



Berkeley Madonna generates the following equations:

```
{ 1: A <--> B }
RXN1 = K1f*A - K1r*B
K1f = 1
K1r = 1
INIT A = 0
INIT B = 0
d/dt(A) = -RXN1
d/dt(B) = +RXN1
```

Now insert the following equation before the chemical reaction object:

```
B = 2.2
```

Choose **Compile Model** from the **Model** or **Compute** menu. Berkeley Madonna returns an Exception error. The problem is that B is defined twice, first by the line you inserted and then by the equations generated by the chemical reaction object. If you enter the above equation after the reaction object, you get a similar error: this time your attempt to define B as a normal equation (converter) is illegal because it has already been defined as a reservoir by the chemical reaction object.

It is possible to override the equations generated from a chemical reaction object. This can be useful if you want to, for example, change the rate of reaction from mass action to some other law. To do this, you need to add a new definition for the rate of reaction variable after the chemical reaction object. It is imperative that you add your overriding definition after the reaction object; otherwise, the object's definition will override your definition, which is probably not what you had in mind.

For example, adding the following line after the chemical reaction object shown above changes the rate of reaction to the following:

```
RXN1 = 0.1
```

Note that this reaction rate is arbitrary and non-physical.

Curve Fitter

Berkeley Madonna can automatically find the values of one or more parameters in your model that minimize the deviation between your model's output and a dataset. Before you use this feature, you should be able to run your model and be confident that it is working properly. You should already have determined integration parameters that give accurate results while minimizing execution time.

To perform a curve fit, choose **Curve Fit...** from the **Parameters** menu. You will see a fairly involved dialog box. Here's what you do with it:

1. Specify which variable in your model you are trying to fit to the external data.
2. Specify the dataset you want to fit the variable to. If you haven't imported the dataset yet, you can do so by clicking the **Import Dataset** button. See Importing Datasets on page 16 for details.

3. Specify one or more parameters in your model which you want to solve for. To do this, choose a parameter from the available list on the left and click **Add**. If you change your mind, click **Remove**.
4. For each parameter you have chosen, you must provide two initial guesses. To set guesses for a parameter, select its name in the parameter list and edit the values displayed in the **Guess #1** and **Guess #2** boxes.
5. You must also provide a minimum and maximum value for each parameter. Both guesses must be within this range. The curve fitter will not set the parameter outside of this range.
6. Specify the fractional tolerance you desire in the solution. For example, if you want the parameter(s) solved to three significant figures, set the tolerance to 0.001. The tolerance must be greater than zero.
7. Once you are satisfied with your setup, click OK to begin the curve fit. Berkeley Madonna will run your model repeatedly until it finds a solution. If you get tired of waiting, click the **Stop** button to abort the fit.
8. When the fit is complete, Berkeley Madonna leaves the parameters set to the values that give the best fit. It then runs your model and plots the fit variable and dataset in a graph window.

While the curve fit is in progress, Berkeley Madonna displays the RMS deviation between the data and best run so far. The deviation is root mean square of the differences between individual data points in the dataset and the corresponding points in the run.

When the curve fit operation complete, Berkeley Madonna will report Root Mean Square (RMS) deviation of the best fit in a dialog box on the graph.

Multiple Fits

By default, the curve fitter fits a single variable in your model to a single imported dataset. However, it can simultaneously fit multiple variables to multiple datasets. To do this, check the **Multiple Fits** box in the **Curve Fit...** dialog and add variable-dataset pairs to the list of fits using the **Add** button. Each variable-dataset pair has a relative weight which allows you to control the importance of each pair. The default weight is one, but it can be changed by selecting the pair and entering a new weight in the box below the list of fits. Berkeley Madonna multiplies the deviation of each pair by its weight, then sums these products to compute the overall deviation to minimize.

Note that when the **Multiple Fits** box is checked, the selections in the **Fit Variable** and **To Dataset** controls have no effect on the curve fit operation; they are used simply to define variable-dataset pairs to add to the **Multiple Fits** list.

Specifying Initial Guesses

There are subtleties involved in specifying initial guesses for each parameter. First, the two guesses must be different.⁸ If one or more parameter's initial guesses do not meet this criterion, the curve fit will stop before the best solution is found.

When you add a parameter to the list, Berkeley Madonna proposes default guesses. Guess #1 is set to 0.5 times the parameter's value. Guess #2 is set to 1.5 times the parameter's value. However, if the parameter's value is zero, the guesses are set to -1 and +1.

Floating-point Exceptions

Some models are very touchy when you vary their parameters. That is, they will not run at all! When this happens, you get a floating-point exception. As the curve fitter runs, it may change the value of a parameter in such a way that the model won't run. The curve fitter can handle this case by backing off, but when it first starts up it must be able to run your model with any combination of your initial guesses. If your model won't run with one of these combinations, the curve fitter won't be able to start. You'll know when this happens if you get an error while the curve fitter is in the initialization phase.

Optimizer

The optimizer feature automatically searches for parameter values that minimize an arbitrary expression. To perform an optimization, choose **Optimize** from the **Parameters** menu. The dialog is set up similar to the Curve Fitter (described on page 51) except that instead you specify an expression to minimize instead of a variable and dataset.

The expression to minimize is evaluated once after each run when variables have reached their final values (TIME = STOPTIME). It should refer to at least one variable in your model. For example, if your model has a reservoir named "Q" representing charge and you want to minimize the final value of this charge, you would simply enter "Q" in the **Minimize Expression** box.

If you want to minimize something other than the final value of a variable, you'll need to define an auxiliary variable in your model that computes the desired quantity at the final time. For example, say you have a variable "temp" representing temperature and you want to adjust several parameters to get the average temperature over time equal to 20 degrees. To do this, you need to define a variable that computes the average temperature. This can be done using a couple of difference equations:

```
INIT sum_temp = 0
NEXT sum_temp = sum_temp + temp
```

⁸ More specifically, the fractional difference between the two guesses should be much greater than the tolerance you specified, i.e.:

$$\frac{|g_1 - g_2|}{\frac{1}{2}(|g_1| + |g_2|)} \gg tol'$$

where g_1 is Guess #1, g_2 is Guess #2, and tol is the desired tolerance of the solution.

```
INIT n = 0
NEXT n = n+1
```

At the final time, the variable `sum_temp` will contain the sum of temperature values for all steps and `n` will contain the number of steps. So, to get the average temperature to 20 degrees, you'd enter the following into the **Minimize Expression** box:

```
ABS(sum_temp/n - 20)
```

Note the use of the absolute value function which ensures that this expression takes on its minimum possible value (zero) when the average temperature is exactly 20 degrees.

Boundary Value Solver

Berkeley Madonna can solve boundary value problems by varying one or more parameters such that specified boundary conditions are met. The requirements for using the boundary value solver are similar to those for the curve fitter. However, instead of using imported datasets, the boundary value solver uses one or more boundary conditions that you specify in the **Boundary Value ODE** dialog box. Berkeley Madonna employs a *multi-dimensional Newton-Raphson* algorithm to identify the optimal parameters.

To solve a boundary value problem, choose **Boundary Value ODE** from the **Modules** submenu in the **Model** menu. Then, complete the dialog box as follows:

1. Specify a boundary condition by: (1) choosing a variable from the **Set** control, (2) entering the target value in the = field, (3) entering the time at which the target value should be reached in the at **X** = field. Then, click the **Add** button. The boundary condition will be displayed in the form `var(x) = y` in the **Boundary Conditions** list on the right. Repeat this procedure for each boundary condition.
2. To remove a condition you don't want, select it from the list and click **Remove**. To change a condition, remove it, edit the appropriate fields, then add it back again.
3. Specify a parameter to solve for by choosing one from the **Available Parameters** list and clicking the **Add** button. The parameter appears in the **Unknowns** list on the right.
4. Adjust the minimum and maximum values for the parameter so that they are just wide enough to cover the range where you expect the solution to be found. Berkeley Madonna will not search for solutions outside of these limits. If the limits are too large, it will take much longer to find a solution or may not be able to find one at all.
5. Repeat steps 3 and 4 for each parameter you want to solve for. Note that the number of parameters must be the same as the number of boundary conditions specified in step 1.
6. To remove a parameter you don't want, select it from the **Unknowns** list and click **Remove**.
7. Specify the relative accuracy you desire in the **Tolerance** field. The default is 0.001 (one part per thousand).

Once everything is set up, start the solver by clicking OK. The solver will run until it finds a solution that meets all of the specified boundary conditions. If a solution cannot be found, the solver will run indefinitely. You can stop it at any time by clicking the **Stop** button.

When the solver finds a solution, it leaves the parameters set to the values that satisfy the boundary conditions and runs your model, plotting the variables used in the boundary conditions. If you stop the solver before it finds a solution, it leaves the parameters set to the values that came closest to the solution and does not run your model.

Sometimes the solver will finish, but specified boundary conditions are not met. This can happen because the tolerance is not stringent enough. For instance, if your model doesn't respond smoothly to small parameter changes close to the solution, the solver may exit without obtaining an adequate solution. To avoid this problem, use a lower tolerance in the **Boundary Value ODE** dialog, or use a smaller stepsize to improve the stability of your model's output.

Sensitivity

The sensitivity feature computes the sensitivity of variables in your model to changes to one or more parameters. The sensitivity of a variable V to a parameter P is computed as follows:

1. Run the model with all parameters at their specified values. We'll call the results of this run $V1(t)$.
2. Adjust parameter P slightly by adding an amount Δ equal to $0.001 * P$. Run the model again, calling the results $V2(t)$. Note that if P is zero, Δ is set to 0.001.
3. Compute the sensitivity $S(t)$ by the following formula:

$$S(t) = (V2(t) - V1(t)) / \Delta$$

When you choose **Sensitivity** from the **Parameters** menu, you can select more than one parameter. Berkeley Madonna will perform one normal run followed by runs where each parameter is perturbed one at a time. It then computes the sensitivity and stores the results as separate runs, one for each parameter you chose. The legend shows the parameter for which sensitivity was measured.

Sensitivity cannot be determined when using variable stepsize algorithms; in order to subtract results of one run from another, Berkeley Madonna needs the step times to match exactly which can only be achieved using fixed stepsize algorithms.

Flowchart Reference

This section describes advanced techniques for working with visual models. It also describes the various commands and preferences for flowchart windows.

Aliases

As your models grow more complex, you'll encounter situations where you want to make a connection (arc or flow) between two icons that are spaced far apart on the flowchart. If you make the connection between these icons, you'll end up with a long arc or flow line that crosses over a bunch of unrelated icons and connections. Not only does this look messy, but it makes it harder to see the essential structure of your model.

One solution to this problem is to move one of the icons you want to connect closer to the other. However, if both icons already have a lot of connections, moving either one will make your flowchart even more disorganized.


A better solution to this problem is to use an alias. Instead of moving one icon closer to the other, you can create an alias for one of the icons and place it next to the other icon. Then you connect the alias to the other icon with a short, simple connection.

An alias is an icon that provides another way to access an existing “original” icon. The most important thing to remember is that an alias **always** refers to an original. You can locate the original by selecting an alias and choosing **Find Original** from the **Flowchart** menu. Aliases look like hollow versions of their originals.

To create an alias, select one or more existing icons and click the **alias** button in the toolbar. Berkeley Madonna will create one alias for each icon you selected. However, note that aliases will only be created for reservoir, flow, and formula icons. You can create as many aliases for a particular original as you want by repeatedly clicking the alias button. These aliases can be placed anywhere in your flowchart.


An alias always has the same name as the original to which it refers. If you change the name of an alias, it changes the name of the original and **all** of its aliases. If you double-click an alias, it opens the icon dialog for the original.

If you make connections to an alias, the effect is identical to that resulting from making the same connections to the original. For example, suppose you have an original formula named X and one arc going from another formula A to X. Also, suppose that you have an arc going from another formula B to an alias of X, like this:



In this case, formula X depends on formula A and formula B, just as if an arc had been drawn from formula B to the original formula X.

When you delete an alias, any arcs connected to it are also deleted. And if it is a reservoir, any flows connected to it are detached and replaced by infinite source/sinks. However, it is possible to delete an alias while preserving its connections. This is called “alias merging” and is accomplished by dragging the alias you want to remove to its original (or another one of its aliases) and dropping it. Upon dropping the alias, it disappears and its connections are merged into the icon on which it was dropped. For example, in the above flowchart you can drag the alias X to the original X and drop it. Berkeley Madonna deletes the alias and connects the arc from formula B to the original X, like this:




Alias merging is very useful when working with submodels as you will soon see.

Submodels

You've already seen how aliases can help make complex models easier to construct and maintain. But as your models continue to grow, you'll soon realize that there simply isn't enough room in a single flowchart window to contain all of your model's elements. And even if you were able to cram your entire model into one full-screen window, it would be difficult to see its overall structure through the forest of arcs, flow lines, and icons.

Berkeley Madonna allows you to break your model down into functional parts by placing them in submodels just as you would place related groups of files into a folder or directory on your computer. Submodels appear as icons in your flowchart. The flowchart containing a submodel icon is referred to as the submodel's parent flowchart. When you double-click a submodel icon, a separate flowchart window is opened showing the contents of that submodel.

Creating and Deleting Submodels

To create a new, empty submodel, click the submodel tool  on the toolbar, drag the mouse over the flowchart and release it. Berkeley Madonna places a new submodel icon on your flowchart.

Often, you'll want to create a new submodel containing some icons that already exist on your flowchart. To do this, select the icons and choose **Group** from the **Flowchart** menu. Berkeley Madonna creates a submodel icon on your flowchart and places the previously selected icons inside that submodel. Note that when icons are moved into a submodel, aliases may be created in order to preserve the connections that existed before the submodel was created. See *Moving Icons Between Submodels* on page 59 for details about how this works.

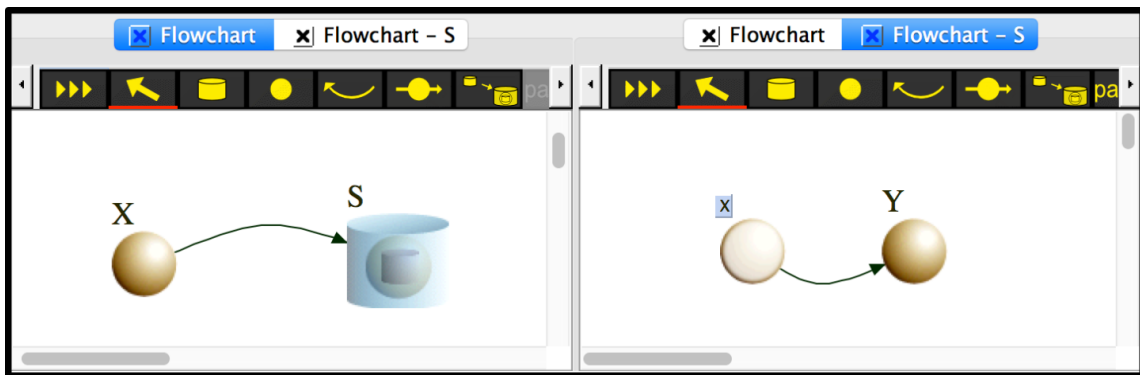
To edit the contents of a submodel, you must first open its flowchart window by double-clicking the submodel's icon. Once opened, you can edit the submodel's flowchart in the same way you edit your model's top-level flowchart.

When you delete a submodel icon, you delete not only the icon but all of the icons it contains. However, you can delete a submodel icon while preserving the icons and connections it contains by selecting one or more submodel icons and choosing **Ungroup** from the **Flowchart** menu. Berkeley Madonna removes the submodel icon and moves all the icons it contained into the parent flowchart.

Making Connections Between Submodels

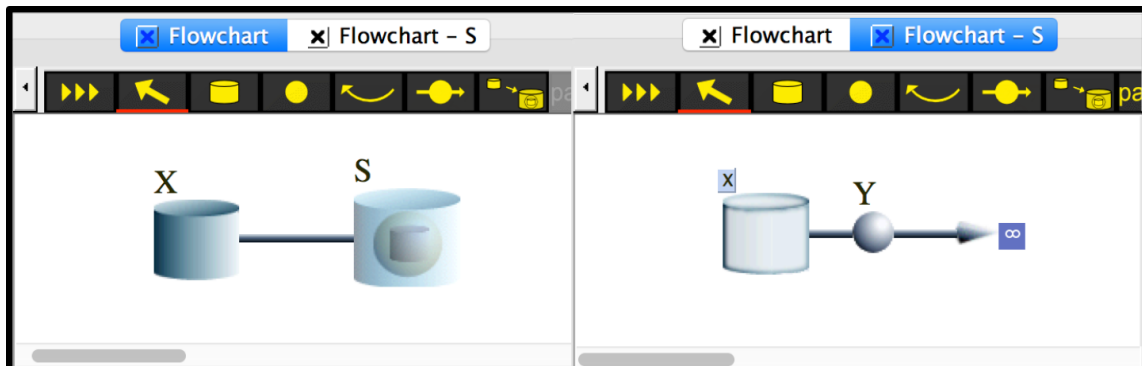
Submodels wouldn't be very useful if you couldn't make connections between icons in different submodels. But simply creating an arc or flow between flowchart windows doesn't work. For example, place a formula icon named X and a submodel icon named S in an empty flowchart window. Open the submodel's flowchart by double-clicking its icon and place a formula icon named Y in its flowchart. Now place the two flowchart windows side-by-side and try creating an arc from X to Y. As you can see, Berkeley Madonna doesn't let you create a connection across two flowchart windows.

So, how do you make a connection between two icons in different submodels? You use an alias. Continuing the above example, create an alias to the X icon and drag it to the submodel icon. This moves the alias from the top-level flowchart into submodel S. Now you can make the connection from X's alias to icon Y in the submodel's flowchart window. Try it. Your flowcharts should end up looking like this:



Notice that when you created the arc from X's alias to Y, another arc was magically created in the top-level flowchart going from X to submodel S. Berkeley Madonna does this to remind you that something in submodel S (in this case, Y) depends on X. This "magic" arc is referred to as a "bundle" since it can represent more than one dependency between X and icons within S. For example, if you create icons A, B, and C within S and connect arcs from X's alias to these icons, the top-level flowchart's appearance will not change. The bundle (arc) from X to S now represents all of the connections from X to icons within S (A, B, C, and Y).

Berkeley Madonna also creates bundles representing one or more flows connected to reservoirs. For example, try this: create a reservoir X in the top-level flowchart. Create an alias to X and move it into S. Then create a flow Y within S going from X's alias to an infinite sink. You should end up with something like this:



When you create the flow Y going from X, Berkeley Madonna creates a bundled flow or “pipe” in the top-level model connecting the original X icon and the submodel. This bundle reminds you that something in S is moving material in or out reservoir X. You could create additional flows in S connected to X and the appearance of the top-level flowchart would not change.

Berkeley Madonna uses a bunch of complex algorithms to ensure that these bundles represent actual flow and arc connections you’ve made. Explaining these algorithms would be difficult and probably not very helpful. There are a couple of simple rules to remember that will help you avoid confusion:

- Bundles are **always** connected between a reservoir, flow, or formula icon and a submodel icon. Whenever you see an arc connected to or from a submodel icon, it is a bundle created automatically by Berkeley Madonna. The same holds for flow bundles (pipes), but they are easier to distinguish since they have no flow icon or label associated with them.
- You cannot directly create or delete bundles. For example, you can’t draw an arc from a submodel to another icon. The only way to delete bundles is to remove the underlying connection(s) they represent.

One final note about bundles. If Berkeley Madonna needs to create a bundle from a submodel icon to another icon that has no representation in the same flowchart window, it will create an alias to that icon. This automatic alias creation is harmless but can be distracting. To minimize the need for additional aliases, try to keep your connections between neighboring submodel levels (from parent to child and vice-versa) as we did in the above examples.

Moving Icons Between Submodels

You’ve already seen that you can create connections between icons in different submodels by creating connections between aliases. But you can achieve the same effect by moving icons from one submodel to another **after** connections to them have already been created. In some ways, this technique is simpler because it enables you to construct parts of your model within a single flowchart window where all icons are accessible at once. Then, you can isolate parts of your flowchart within their own submodels and let Berkeley Madonna worry about preserving connections between them.

To illustrate how this method of construction differs, let’s try to recreate the example model we discussed earlier. Create a new flowchart and place two formula icons named X and Y. Draw an arc from X to Y. Now, you decide that you want to place Y in its own submodel. To do this, select Y and choose **Group** from the **Flowchart** menu. Now Y has been replaced by a submodel icon named “S1”. Open S1 and you’ll see it has the same structure as when you created the alias to X yourself. If you want to make things look exactly like they did before, you can change the name of the submodel to “S”.

Once you’ve created a submodel or two, you may want to move icons from one submodel to another. Berkeley Madonna supports this kind of rearrangement in two ways. First, you can move a set of icons from a flowchart into one of its submodels (children) by dragging the icons to the submodel icon and dropping them inside. Second, you can move a set of icons from a submodel into its parent flowchart (the flowchart containing that submodel’s icon) by selecting the icons and choosing **Move To Parent** from the **Flowchart** menu.







As you begin to move icons around between submodels, you’ll find that Berkeley Madonna can create a lot of aliases along the way. Or you might find some icons moving that you didn’t want moved. Here are some tips that will help you understand what’s going on:








- When you drag an icon from a parent model into a child, its connections to other icons are also moved. For example, say you have a reservoir connected to a flow and a formula in your top model. If you drag this reservoir to a submodel, the arcs and pipes that connect this reservoir to the flow and formula are also moved to the submodel. After the connections are moved, aliases are created in the submodel as needed so that the other ends of these connections are maintained.
- When you move icons from a submodel into its parent, Berkeley Madonna will create aliases in the parent model to preserve any connections that were moved. This behavior can lead to a proliferation of aliases in the parent model. These aliases can often be moved back into the submodel and then merged with their originals to reduce clutter.
- When deciding where to place original icons in your model, remember that it's best to put originals high enough in the hierarchy so that all their references come from other icons in the same flowchart (including submodels). Try to avoid putting icons down in a submodel when they are referenced by other icons in the parent model or sibling submodels.




Take a look at the "HH Axon (Submodels)" model in the "Examples" folder. Although it is somewhat involved, each submodel is easy to understand and you can get a good view of the model's overall structure just by looking at the top-level flowchart.

Flowchart Toolbar

The flowchart toolbar appears at the top of every flowchart window in Berkeley Madonna. The active tool is indicated by a red bar drawn beneath it. When you click the mouse in the flowchart, the function performed depends on the active tool as described in the table below. After using most tools, the **Select** tool is activated. However, you can make certain tools remain active after use, or "sticky", by holding down the control key when selecting the tool. When the active tool is sticky, it is denoted by a thicker red bar.

Tool	Description
Select 	Use the select tool to select one or more icons. Clicking on an icon selects it and deselects any previously-selected icons. Shift-clicking an icon toggles its selection status without deselecting previously selected icons. You can also click and drag the mouse icon to select a group of icons. To deselect all icons, click an empty space in the flowchart. Normally, you don't need to choose this tool since it is activated after using any of the other tools. However, if you use the "sticky tool" feature described above, you'll need to explicitly activate the Select tool when you're done using the sticky tool.
Reservoir 	Use the reservoir tool to create a new reservoir on the flowchart. There are two ways to do this. One is to select the tool by clicking once, then moving over the flowchart and clicking again. The other way is to click the reservoir tool, drag the mouse over the flowchart, then release the mouse. If you want to create many reservoirs rapidly, control-click the reservoir tool to make it sticky, then repeatedly click in the flowchart. Remember to click the select tool when you're done using the sticky reservoir tool.
Formula 	Use the formula tool to place formula icons on the flowchart. You can use either the two-step, one-step (drag), or sticky methods as described for reservoir tool.
Arc 	Use the arc tool to create a dependency relationship between two icons. After selecting the tool, move the mouse to the source icon, click and drag the mouse to the destination icon, and release the mouse. This creates an arc pointing from the source to the destination icon representing the destination icon's dependency on the source icon. The source/destination icons must be reservoirs, formulas, or flows. You can use the sticky tool (control click) feature if you need to create many arcs quickly.
Flow 	Use the flow tool to represent a flow of material between a reservoir and an infinite source/sink or between two reservoirs. After selecting the tool, click the mouse on the source, then drag the mouse to the destination and release it. If you want the source or destination to be an infinite source or sink, respectively, start or end the drag at an empty area of the flowchart and Berkeley Madonna will create the source/sink for you. If a source/sink is created but you intended to connect that end to a reservoir, drag the source/sink to the reservoir to make the connection. You can use the sticky tool feature to create a number of flow connections quickly.
Submodel 	Use the submodel tool to organize portions of your model into a hierarchical structure. This feature is described in Submodels on page 57.

Tool	Description
Text 	Use the text tool to create text labels on the flowchart. After selecting the tool, click on the flowchart to place an empty text box. You can then type in your label text. Unlike other tools, this tool is always sticky. When this tool is active, you can click on any text label (including those associated with icons) to position the insertion point for editing.
Align 	Use the align tool to align the top, bottom, left, or right edges of the selected icons. The direction of alignment depends on where you position the mouse within the tool. Watch the help text as you move the mouse within this tool.
Grid 	Use the grid tool to turn the grid on and off. When the grid is on and you move an icon, its position is adjusted so that the edge closest to a gridline coincides with that gridline.
Alias 	Use the alias tool to create an alias for the selected icon(s). An alias is a separate representation of the original icon. Alias icons appear as hollow versions of their originals. See Aliases on page 55 for more information.
Show/ Hide Formulas 	Use the show/hide formulas tool to toggle the visibility of formulas (and their connecting arcs) in the flowchart. Hiding formulas and their arcs can make a complex flowchart easier to understand. This is equivalent to toggling the Formulas checkbox in the Icon Visibility dialog.
Parent 	Use the parent tool to move from within a current submodel to the parent flowchart one level above. In the upper most model, the parent tool will be greyed out.
Delete 	Use the delete tool to remove a graphical element from your model. Select an element then click the delete tool to remove it. All connections (arcs) to that element will also be deleted. If you have highlighted several elements from your model, they will all be deleted when you press delete. If you delete a submodel, all contents of the submodel will also be deleted.

Tool	Description
Global	The global tool will open the Globals Window for you if it is not already open in your model. If the Globals Window is already open, it will do nothing. In this latter case, click on the Globals button at the top of the window to move to this page.
	
Multiline Text	Use the multiline text tool to create text labels on the flowchart. After selecting the tool, click on the flowchart to place an empty text box. You can then type in your label text. When this tool is active, you can click on multiline text label (including those associated with icons) to position the insertion point for editing.
	
Run	The run tool will run the current model and take you to a Graph Window. If one is not open, it will create one for you.
	

Flowchart Commands

Commands in the **Flowchart** menu apply to the active flowchart window. In addition, there are a few other flowchart-related commands in other menus. Commands shown in the following table appear in the **Flowchart** menu unless otherwise specified.

Command	Description
New Flowchart Document (File menu)	Creates a new visual model and opens an empty top-level flowchart window.
Cut (Edit menu)	Copies the selected icons to the clipboard and removes them from the flowchart.
Copy (Edit menu)	Copies the selected icons to the clipboard.
Paste (Edit menu)	Pastes the icons on the clipboard to the active flowchart. If the names of any pasted icons clash with existing icons, they are given new names.
Select All (Edit menu)	Selects all icons in the active flowchart.
Show Icon (Edit menu)	Selects the icon corresponding to the specified equation in the equation window. The specified equation is indicated by the position of the insertion point or the beginning of the selected text. Available only when the equation window is active.
Copy Image	Copies the entire flowchart to the clipboard as a picture. This picture can then be pasted into other applications.
Icon Info...	Displays the icon dialog for the selected icon. For submodel icons, this command opens the submodel's flowchart window. For text labels

Command	Description
	created with the text tool, this command displays the text dialog which allows you to change the text style. This command is available only when a single icon is selected and has the same effect as double-clicking an icon.
Find Original	Locates and selects the original icon for the selected alias icon. Available only when a single alias icon is selected.
Group	Places the selected icons into a submodel. See Submodels on page 57 for more information.
Ungroup	Moves icons from the selected submodel into its parent model, then deletes the submodel. Select the submodel prior to clicking ungroup – not the contents of the submodel. See Submodels on page 57 for more information.
Move To Parent	Moves the selected icons to the parent model. See Submodels on page 57 for more information.
Show Parent Window	Activates the flowchart window containing the active flowchart's submodel icon. See Submodels on page 57 for more information.
Redraw	Redraws the flowchart window. Useful when a portion of the flowchart appears corrupted.
Background Color...	Displays a dialog which enables you to change the flowchart's background color.
Import Image...	Replaces the <i>selected</i> icon(s) default image with a custom image from an external file. The image file must be in GIF, PNG, or JPEG format. If no icon is selected, the background will be set to the image.
Remove Image	Removes the custom image from the selected icon(s) or background. This command is available only when one or more icons with custom images are selected. If nothing is selected, it will remove the custom background image.
Icon Visibility...	Displays a dialog which enables you to hide different kinds of icons and connectors in the flowchart.
Tools	This submenu provides another way to select various tools in the toolbar.
Arrange » Front	Makes the selected icons appear in front of all other icons.
Arrange » Back	Makes the selected icons appear behind all other icons.
Arrange » Align	Aligns the top, bottom, left, or right edges of the selected icons. Same as the align tool on the toolbar.
Arrange » Distribute	Distributes the selected icons evenly across in the flowchart in the horizontal or vertical direction.
Arrange » Snap To	Aligns the selected icons with the gridlines. Works even when the grid is

Command	Description
Grid	not visible.
Show Flowchart	Shows the active model's hidden flowchart windows. If none of the flowchart windows are hidden, it simply opens the top-level flowchart. This command changes to Hide Flowchart when at least one flowchart window is open and none of the flowchart windows are hidden.
Hide Flowchart	Hides the active model's flowchart windows. The hidden windows can be quickly restored by choosing Show Flowchart . Note that closing a flowchart window by clicking its close box closes the window rather than hiding it. This distinction is important because only hidden flowchart windows are restored when Show Flowchart is chosen.
Discard Flowchart	Removes the flowchart from the active model, converting it into a plain text model. Once this is done, the flowchart cannot be recovered. However, you can continue to edit the equations in the equation window.

Flowchart Preferences

The behavior of Berkeley Madonna's flowchart editor can be customized using the **Flowchart** page of the **Preferences** dialog as follows:

Preference	Description
Arc Style	Specifies the type of curve used for dependency arcs between icons. The default is a "single point" Bezier curve whose trajectory is adjusted via a single control point. "Double point" bezier curves provide more control over the trajectory but require twice as much fiddling to position since they have two control points. The "straight line" option makes all arcs into straight lines.
Grid On Initially	If checked, the grid in flowchart windows is turned on when they are first opened.
Show Drag Alerts	When checked (the default), error messages are displayed if you attempt to make an invalid connection with a flow or arc, for example, trying to draw an arc from one reservoir to another. Turning off this preference prevents these error messages from being displayed.
Grid Spacing	Specifies the distance between gridlines in the flowchart.
Max. Text-to-Owner Distance	Specifies the maximum distance between and icon and its text label. Berkeley Madonna ensures that the distance between the icon's enclosing rectangle and its text label is not greater than this distance.

Customizing Icons

The **Import Image** command discussed earlier can be used to substitute a custom image (in a GIF, PNG, or JPEG file) for Berkeley Madonna's default icon images or background. This mechanism has some drawbacks:

- It is difficult to tell when icon with a custom image has been selected because it doesn't change appearance.
- Custom icons considerably increase the size of the model file. This is especially wasteful when the same images are used in many models.

Berkeley Madonna provides an alternative way to customize the images used for all instances of a particular type of icon (GIF only). This method makes it possible to change the entire look of your flowcharts. To customize icons first *create and select a folder* with the icons:

- 1) Put the icons in any directory you like.
- 2) In **Preferences...** go to the **Flowchart** tab and select **Custom Icons**. If you can't see the Custom Icons button, expand the window downward. Then use the **Choose** button to select any single file in the directory with the custom icons.
- 3) Hit **OK** in the Preferences window and exit Berkeley Madonna.

Next time you start Berkeley Madonna, any images in the selected directory with names matching the flowchart icon names will be used. If any icons are not provided, the default icons will be used instead. To revert to the old icons, either deselect "Custom Icons" or move the custom icons out of the custom icon directory. Now *name your images* with the following convention:

`<type><select><alias><size>.gif`

where:

- `<type>` is a three-letter code indicating the type of icon:

con	formula icon (converter)
flo	flow icon
mod	submodel icon
res	reservoir icon
- `<select>` is a single letter indicating whether the image is for a selected icon or a normal (unselected) icon:

n	Normal icon (not selected)
s	Selected icon
- `<alias>` is a single letter indicating whether the images is for an alias icon or a normal (non-alias) icon:

n	Normal icon (not an alias)
a	Alias icon
- `<size>` is a single letter denoting the size of the icon:

s	Small icon
l	Large icon

Here are some examples:

resnas.gif Reservoir alias icon, not selected, small size
modsnl.gif Submodel icon, selected, large size

Note that since there is no such thing as an alias to a submodel icon, images named **modXaX.gif** are not used.

The customized icon images will be recognized the next time Berkeley Madonna is launched.

Known Bugs and Limitations

1. The flowchart editor cannot generate equations employing some of Berkeley Madonna's more advanced equation features, but they can be used in a visual model by typing them into the Globals Window. These include:
 - Array equations.
 - Upper and lower bounds (LIMIT statement).
 - Built-in root finder equations.
 - Renaming system symbols (RENAME statement).
 - Exposing a limited subset of a model's symbols (DISPLAY statement).

Also, the Chemical Reactions feature cannot be used with visual models.

We may add support for some of these features in a future release of Berkeley Madonna.

2. Conveyors, ovens, and queues cannot be represented with stock icons and flows like reservoirs. Instead, they must be defined using formula icons.
3. The flowchart icon dialogs allow you to define graphical functions using Berkeley Madonna's GRAPH() built-in function. However, the Graph Input dialog is fairly buggy and difficult to use. This feature will be improved in version 10.
4. Several features from Berkeley Madonna version 8 are not present in version 9-10.4, and their deprecation has been noted throughout including: Oscilloscope mode, Check DT.