

Manifold 0.9 User Guide


October 8, 2012

1 Introduction

Manifold is a parallel discrete event simulation framework for simulation of modern multicore computer architectures. The software package consists of a parallel simulation kernel, a number of component models, and a few ready-to-use simulator programs that use the component models to build and simulate system models. Users can also port third-party components to Manifold and build system models with them. This user guide describes how to obtain Manifold source code, and how to build and run the simulator programs.

2 Overview

Manifold is designed for parallel simulation of multicore systems. The general simulation system is shown in Figure 1 below.



At run-time, instruction streams are fed to the multicore system model for simulation. Example sources of instructions include PIN [1] trace files and the QSim [2] multicore emulator. Components of the system model can be assigned to different host machines for parallel simulation.

The following are the general steps that a simulator program needs to follow to create a system model for simulation.

- Instantiating the various component models, such as processor model, cache model and so on.
- Connecting the components with Manifold links.
- Setting a simulation stop time.
- Starting the simulation by calling `Manifold :: Run()`.

The simulator programs that are part of the distribution package can serve as examples for how to write simulator programs with Manifold. The component models that are included in the package can be used in building system models. The user can also port third-party components to Manifold and build system models using such components.

Major features of Manifold include the following:

- Supporting both sequential and parallel simulations.
- Supporting three simulation paradigms: discrete time event-driven simulation, time-stepped simulation, and mixed event-driven and time-stepped simulation.
- Supporting the QSim multicore emulator front-end.
- The Energy Introspector (EI) power, energy, thermal, and reliability modeling library for integration with multicore processor models.
- Standard interfaces between components allow mix-and-match of components.
- Open software architecture allows easy porting of third-party components.

3 Current Release

The current release is Release 0.9. The software is distributed as a source code package that contains the following:

- The parallel simulation kernel.
- Models:
 - Zesto: a cycle-level x86 processor model.
 - Mcp-cache: a coherence cache model.
 - Iris: a cycle-level interconnection network model.
 - CaffDRAM: a DRAM controller model.
- Simulator programs.

Zesto

- This model can accept instruction streams from three different sources: PIN trace files, QSim server, and QSim library.
- Both in-order and out-of-order models are included.

Mcp-cache:

- Supports the MESI protocol.

Iris:

- Supports ring and torus.
- Supports virtual channels.
- Supports two virtual networks.
- Supports flow-control between routers, between router and interface, and between interface and terminal.
- Supports single-flit packets.

CaffDRAM:

- Supports flow control between the memory controller model and the network model.

Testing and Portability

- Manifold has been tested on the PARSEC and SPLASH-2 benchmarks.
- Manifold has been tested under the the following operating systems:
 - Ubuntu 10.04, 12.04.
 - Fedora release 17.
- Manifold has been tested with Openmpi and MPICH2.

4 Source Code Directory Structure

The Manifold source code is organized as follows:

```
ROOT
|... code
|... doc
|... kernel
|... models
|    |... cache
|    |    |... mcp-cache
|    |... memory
|    |    |... CaffDRAM
|    |... network
|    |    |... iris
|    |... processor
|    |    |... zesto
|... simulator
|... smp
|    |... common
|    |... config
|    |... QsimClient
|    |... QsimLib
|    |... TraceProc
```

where *ROOT* represents the root of the source tree.

Under each of *kernel*, *mcp-cache*, *iris*, and *zesto*, there is a subdirectory *doc/doxygen* that contains a user guide in Doxygen format for the respective component.

The simulator directory

The *simulator* directory hosts a total of six simulator programs, two each under *QsimClient*, *QsimLib*, and *TraceProc*. The difference is the source of instruction streams. Programs under *QsimClient* use QSim server to get instructions. Those under *QsimLib* are built with the QSim libraries. And those under *TraceProc* use PIN traces.

The common code of the simulator programs is located in `common`. The directory `config` contains configuration files for the simulator programs.

5 Build Process Overview

To build and run the simulator programs that are part of the software package, you will need to perform the following steps:

- [Optional] Install required packages.
- [Optional] Download and build QSim.
- Build Manifold libraries.
- Build the simulator program(s).
- Run the simulators.

The simulators can respectively take instructions from three different sources: trace files, QSim library, and QSim server. Depending on which source you use, some of the steps above may be optional.

The following explains each step in detail.

6 Install Required Packages

Before you proceed, you need to install the following required packages.

- `mpi`: We have tested with `openmpi` and `mpich2`, so these two packages are recommended.
- `libconfig++`: The simulators require this package.

7 Download and Build QSim

If you choose to use QSim to get instructions, you need to build and install QSim first.

7.1 Download

QSim is available through SVN check out. The latest version at this writing is 0.1.5, at the following address:

- <https://sst-simulator.googlecode.com/svn/qsim/tags/qsim-0.1.5>

7.2 Build and Installation

Instructions for building and installing QSim can be found in the `INSTALL` file in the root directory of QSim source code.

In addition to the QSim libraries, you also need to do the following:

- build the QSim server.
- build and install the QSim client library.

All the instructions are in the `INSTALL` file.

After you are finished, your installation directory should look like the following, assuming `QSIM_INSTALL` is the root of the installation directory.

```
$ ls <QSIM_INSTALL>/lib
libqemu-qsim.so  libqsim-client.so  libqsim.so

$ ls <QSIM_INSTALL>/include
mgzd.h  qsim-client.h  qsim.h  qsim-load.h  qsim-net.h  qsim-regs.h  qsim-vm.h
```

8 Download and Build Manifold Libraries

There are two ways to download Manifold source code: from the Manifold website or through SVN checkout. Depending on which way is used to obtain the source code, the build process is slightly different.

8.1 Download Manifold source package

Manifold source package is available at the Manifold website:

- <http://manifold.gatech.edu/download>

After download, follow the following instructions to build the manifold libraries:

1. Untar the source package.
`$ tar xvfz manifold-0.9.tar.gz`
2. Go to the `code` subdirectory.
`$ cd manifold-0.9/code`
3. Run `configure` and `make`.
`$./configure [--prefix=INSTALL_PATH]`
`$ make`

The default installation directory is `/usr/local/lib`. If you want to install in a different location, the path of that location should be passed to `configure`. In addition, if QSim is installed in a location other than the default, you need to tell `configure` that location. Options that you can specify for `configure` are described below.

4. Optionally, install the libraries.

```
$ make install
```

8.2 Download Manifold source code through SVN checkout

Manifold source code is available through SVN checkout at the following address:

- <https://svn.ece.gatech.edu/repos/Manifold/tags/0.9>

To build the un-packaged source code, you need to have autotools package installed on your machine.

1. From the `code` subdirectory, run `autoreconf`.

```
$ cd code
```

```
$ ./autoreconf -si
```

This would create the `configure` script.

2. Run `configure` and `make`.

```
$ ./configure [--prefix=INSTALL_PATH]
```

```
$ make
```

3. Optionally, install the libraries.

```
$ make install
```

8.3 Configure options

This section describes all of the options you can use when running the `configure` script.

- `--prefix=PREFIX`

By default, the header files and libraries will be installed in `/usr/local/include` and `/usr/local/lib`, respectively. If you want to install the files somewhere else, you should use this option, and the files will be installed in `PREFIX/include` and `PREFIX/lib`, respectively.

- `--disable-para-sim`

By default, the Manifold libraries are built for parallel simulation with MPI. If you do not want to use MPI and therefore only run the simulators in sequential mode, you need to specify this option to disable parallel simulation.

- `--enable-kernel-large-data`

By default, the maximum size of data that are sent between components is 1024 bytes. If this is not big enough, or the maximum size is not known in advance, then this option should be used.

- `--disable-stats`

By default, the Manifold kernel and computer architecture models all collect statistics at run

time. Use this option to disable run-time collection of statistics.

- **KERINC=KERNEL_LOCATION**

This option specifies where the kernel header files are installed. This is useful when the kernel and the models are built separately.

- **QSIMINC=QSIM_LOCATION**

This option specifies the location where QSim is installed. By default, QSim is installed under `/usr/local`. This option is useful when QSim is installed in a different location.

9 Build the Simulator Program

The simulator programs are located in `ROOT/code/simulator/smp`. There are three subdirectories for three types of simulators, based on how they get instructions:

- Programs under `QsimClient` use QSim server to get instructions. To build these simulators, you must first build and install QSim.
- Programs under `QsimLib` use QSim libraries. To build these simulators, you must first build and install QSim.
- Programs under `TraceProc` use trace files in the format of Intel's Pin.

In addition there are two other subdirectories:

- Subdirectory `common` contains code that is shared by all simulators.
- Subdirectory `config` contains configuration files that are shared by all simulators.

To build the simulators, follow the following steps. Here we use the simulators under `QsimClient` as an example.

1. Go to the simulator source directory.

```
$ cd ROOT/code/simulator/smp/QsimClient
```

2. Run `make`. It is likely that you need to modify the Makefile so the header files and libraries can be found.

```
$ make
```

10 Start the Simulator

In each of the three subdirectories, i.e., `QsimClient`, `QsimLib`, and `TraceProc`, there is a program called `smp_11p`. This program simulates the following system model, where each core node has a processor core, a private L1 cache, and a shared L2 slice.

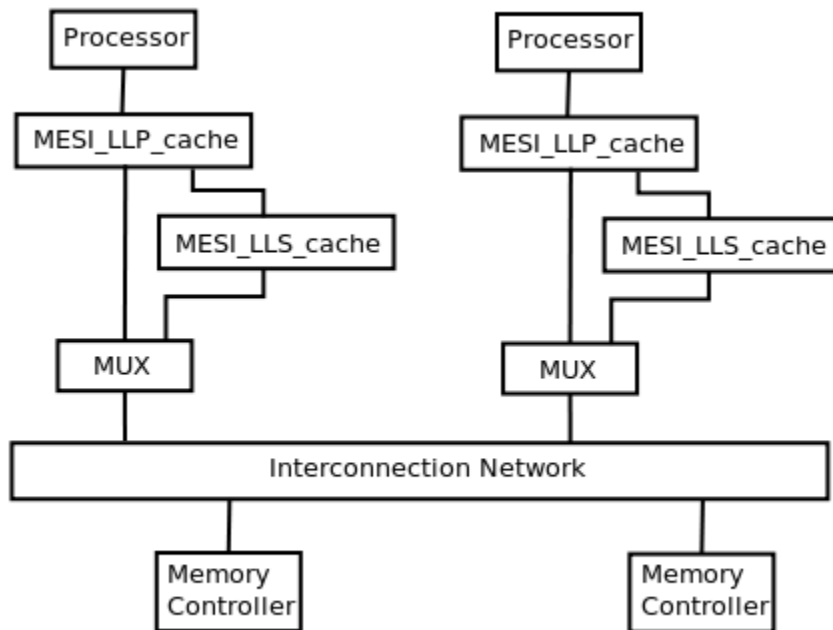


Figure 1 System Model Simulated by smp_llp.

We describe how to start the simulators in each of the three subdirectories.

10.1 Start the Simulators in QSimClient

These simulators require QSim server be started first.

To start the QSim server, run the following commands:

```
$ cd QSIM_ROOT/remote/server
$ make
$ ./server <port> <state_file> <benchmark>
```

where

- `<port>` is the TCP port number the server uses. Use any number you want.
- `<state_file>` is the state file. QSim is an emulator of a multicore shared-memory machine. The state file is the snapshot of the emulated machine after the OS has booted.
- `<benchmark>` is the tar file containing the application program and its data. See QSim instructions on how to build benchmark tar files.

After the QSim server has started, the simulator can be started.

If QSim is installed in `/usr/local`, do the following,

```
$ cd SIMULATOR_ROOT
$ mpirun -np <NP> <prog> <conf_file> <zesto_conf_file> <server> <port>
```

If Qsim is not installed in `/usr/local`, do the following, assuming QSim installation path is `QSIM_INSTALL`.

```
$ cd SIMULATOR_ROOT
$ QSIM_PREFIX=<QSIM_INSTALL> LD_LIBRARY_PATH=<QSIM_INSTALL>/lib mpirun
-np <NP> <prog> <conf_file> <zesto_conf_file> <server> <port>
```

where

- `<NP>` is the number of logical processes (LPs), or MPI ranks. For parallel simulation, currently the simulators support 1, 2, or $N+1$ LPs, where N is the number of simulated cores.
- `<prog>` is the simulator, including `smp_llp`, and `smp_lll2`.
- `<conf_file>` is the configuration file for the system being simulated. The system configuration is defined in `libconfig` format.
- `<zesto_conf_file>` is the configuration file for Zesto, our processor model.
- `<server>` is the name or IP address of the QSim server.
- `<port>` is the TCP port number used by the QSim server.

For example:

```
$ mpirun -np 2 smp_llp ../config/conf2x2_iris_torus_llp.cfg ../config/6.cfg
localhost 12345
```

10.2 Start the Simulators in QSimLib

Simulators in this subdirectory can only be run with 1 LP, or in sequential mode.

If QSim is installed in `/usr/local`, do the following.

```
$ mpirun -np 1 <prog> <conf_file> <zesto_conf_file> <state_file> <benchmark>
```

If Qsim is not installed in `/usr/local`, do the following, assuming QSim installation path is `QSIM_INSTALL`.

```
$ QSIM_PREFIX=<QSIM_INSTALL> LD_LIBRARY_PATH=<QSIM_INSTALL>/lib mpirun  
-np 1 <prog> <conf_file> <zesto_conf_file> <state_file> <benchmark>
```

where

- `<prog>` is the simulator, including `smp_1lp`, and `smp_1112`.
- `<conf_file>` is the configuration file for the system being simulated. The system configuration is defined in `libconfig` format.
- `<zesto_conf_file>` is the configuration file for Zesto, our processor model.
- `<state_file>` is QSim's state file.
- `<benchmark>` is the application tar file.

For example:

```
$ mpirun -np 1 smp_1lp ../config/conf4x1_ring_1lp.cfg ../config/6.cfg myState_16  
myBench.tar
```

10.3 Start the Simulators in TraceProc

These simulators use Pin traces.

```
$ mpirun -np <NP> <prog> <conf_file> <zesto_conf_file> <trace_file_basename>
```

where

- `<NP>` is the number of logical processes (LPs), or MPI ranks. For parallel simulation, currently the simulators support 1, 2, or $N+1$ LPs, where N is the number of simulated cores.
- `<prog>` is the simulator, including `smp_1lp`, and `smp_1112`.
- `<conf_file>` is the configuration file for the system being simulated. The system configuration is defined in libconfig format.
- `<zesto_conf_file>` is the configuration file for Zesto, our processor model.
- `<trace_file_basename>` is the base name of the trace files. All trace files must have the same base name and be named `<base_name>0`, `<base_name>1`, `<base_name>2`, etc. For example, if the trace files are `myFile0`, `myFile1`, then the base name is `myFile`.

For example:

```
$ mpirun -np 2 smp_1lp ../config/conf2x2_iris_torus_1lp.cfg ../config/6.cfg  
myTrace
```

11 Common Problems

The following is a list of commonly encountered problems, and how to solve them.

- `mpirun: command not found`
Solution: If you are using `openmpi`, install the `openmpi-bin` package.
- `simulation_stop` has incorrect type.
Solution: Open the configuration file, append an 'L' to the number you specify for `simulation_stop`. For example, if it was "`simulation_stop = 1000`", change it to "`simulation_stop = 1000L`".
- `cp: cannot stat './libqemu.so': No such file or directory`
`system("cp ./libqemu.so /tmp/qsim_WKlK7m") returned 256.`
Solution: Specify `LD_LIBRARY_PATH` as shown above.
- `cp: cannot stat '/usr/local/lib/libqemu-qsim.so': No such file or directory`
`system("cp /usr/local/lib/libqemu-qsim.so /tmp/qsim_eIwV0x") returned 256.`
Solution: Specify `QSIM_PREFIX` as shown above.

12 References

- [1] www.pintool.org, Pin - A Dynamic Binary Instrumentation Tool, www.pintool.org.
- [2] www.cdkersey.com, QSim - QEMU for Simulators, www.cdkersey.com/qsim-web.