

Flexus Getting Started Guide

Flexus 4.0.0

06/15/2010

Guide authors: Pejman Lotfi-Kamran, Evangelos Vlachos

SimFlex Project Team

Mohammad Alisafae, Eric Chung, Michael Ferdman, Brian Gold,
Jangwoo Kim, Pejman Lotfi-Kamran, Onur Kocberber, Ippokratis Pandis, Djordje Jevdjic,
Jared Smolens, Stephen Somogyi, Evangelos Vlachos, Stavros Volos, Jason Zebchuk,
Babak Falsafi, Nikos Hardavellas and Tom Wenisch



**Computer Architecture Lab at
Carnegie Mellon**

Legal Notices

Redistributions of any form whatsoever must retain and/or include the following acknowledgment, notices and disclaimer:

This product includes software developed by Carnegie Mellon University.

Copyright 2010 by Mohammad Alisafae, Eric Chung, Michael Ferdman, Brian Gold, Jangwoo Kim, Pejman Lotfi-Kamran, Onur Kocberber, Ippokratis Pandis, Djordje Jevdjic, Jared Smolens, Stephen Somogyi, Evangelos Vlachos, Stavros Volos, Jason Zebchuk, Babak Falsafi, Nikos Hardavellas and Tom Wensch for the SimFlex Project, Computer Architecture Lab at Carnegie Mellon, Carnegie Mellon University.

For more information, see the SimFlex project website at:

<http://parsa.epfl.ch/simflex/>

You may not use the name "Carnegie Mellon University" or derivations thereof to endorse or promote products derived from this software.

If you modify the software you must place a notice on or within any modified version provided or made available to any third party stating that you have modified the software. The notice shall include at least your name, address, phone number, email address and the date and purpose of the modification.

THE SOFTWARE IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY OF ANY KIND, EITHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE SOFTWARE WILL CONFORM TO SPECIFICATIONS OR BE ERROR-FREE AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, OR NON-INFRINGEMENT. IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES, ARISING OUT OF, RESULTING FROM, OR IN ANY WAY CONNECTED WITH THIS SOFTWARE (WHETHER OR NOT BASED UPON WARRANTY, CONTRACT, TORT OR OTHERWISE).

Introduction

Computer architects have long relied on software simulation to evaluate the functionality and performance of architectural innovations. Unfortunately, there are many hurdles which make simulation a design bottleneck. First, modern cycle-accurate simulators are several orders of magnitude slower than real hardware. Second, the growing levels of integration continue to result in commensurate increases in computer system size and complexity. Third, benchmarking commercial server software requires full-system simulation including the simulation of peripheral devices and OS code. Finally, because conventional simulators are optimized for speed, they are often not component-based and as such are inflexible and difficult to modify.

The SimFlex project has developed simulation tools and measurement methodology to enable fast, accurate, and flexible performance evaluation of uni- and multiprocessor systems running unmodified commercial applications. SimFlex combines rigorous statistical sampling, to choose application subsets for measurement, with reusable checkpoints of system state, to enable rapid simulation of the selected measurement. Together, these techniques enable ~10,000 times reduction in simulation time relative to cycle-accurate simulation without sampling and up to 1000-way simulation parallelism over a cluster of simulation hosts.

Flexus

Flexus is a family of component-based C++ computer architecture simulators that build on Virtutech Simics's Micro-Architecture Interface (MAI) to enable full-system timing-accurate simulation of uni- and multiprocessor systems running unmodified commercial applications and operating systems.

Flexus encompasses both a simulation infrastructure and default simulation models. A simulator is composed of individual modules that are hooked together during compilation. A module is often the equivalent of a single hardware structure—for example, a branch predictor or a cache. A key strength of Flexus is its isolation of components: one implementation of a particular module can be swapped for a different implementation without requiring changes to any other modules. This flexibility also allows a particular simulator to be tailored to the needs of a specific research hypothesis. If memory system performance is being evaluated, a simple bandwidth-based processor pipeline might be sufficient. Conversely, a study that closely examines microarchitecture could use a simple memory system model. The Flexus core provides services, such as scheduling and statistics, which are common and useful to all simulators.

Virtutech Simics enables full-system simulation. Simics is a functional simulator that allows unmodified commercial operating systems and applications to boot and run. Flexus can hook into Simics and see the instruction stream that a real system would execute. Flexus can also control Simics's timing, so as to model out-of-order effects and speculative techniques.

Flexus is designed to support the simulation sampling and checkpointing methodologies developed by the SimFlex research project. The keys to this support are *flex-points*, checkpoints that store the snapshots of the state of Flexus components alongside Simics checkpoints of programmer-visible state, and *stat-manager*, a tool for processing Flexus statistics output to create reports and compute statistical confidence in results.

Flexus's component-based design enables easy creation of several components that all model the same hardware at various levels of timing fidelity. These components all share the same format for storing state in flex-points. In this way, a simple, fast simulator can rapidly construct a flex-point library, which can then be measured using a more detailed detailed simulator. For example, the Flexus `CMP.L2Shared.Trace` simulator creates flex-points that can be used for detailed out-of-order timing simulation with `CMP.L2SharedNUCA.000`.

The Flexus Getting Started Guide assumes that readers have some familiarity with the SimFlex experiment methodology and the SMARTS and Live-points simulation sampling techniques. If you have not already done so, we strongly recommend reviewing the SimFlex tutorial slides, and reading the following publications (all available on the SimFlex web site at <http://parsa.epfl.ch/simflex/>):

- T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, J. C. Hoe, "SimFlex: Statistical Sampling of Computer Architecture Simulation," *IEEE Micro Special Issue on Computer Architecture Simulation*, July/August 2006.

- T. F. Wenisch, R. E. Wunderlich, B. Falsafi, J. C. Hoe, "Simulation Sampling with Live-points," *Proceedings of the International Symposium on Performance Analysis of Software and Systems (ISPASS)*, March 2006.
- R. E. Wunderlich, T. F. Wenisch, B. Falsafi, J. C. Hoe, "SMARTS: Accelerating Microarchitecture Simulation via Rigorous Statistical Sampling," *Proceedings of the 30th International Symposium on Computer Architecture (ISCA)*, June 2003.

This release of Flexus includes several example Flexus simulation models:

- `UP.Trace` A non-timing uni-processor simulator
- `UP.OoO` Out-of-order uni-processor simulator
- `UP.Inorder` In-order uni-processor simulator
- `CMP.L2Shared.Trace` A non-timing chip-multiprocessor simulator with shared L2 to study cache miss rates and to construct flex-points for `CMP.L2SharedNUCA.Inorder` and `CMP.L2SharedNUCA.OoO`
- `CMP.L2SharedNUCA.Inorder` In-order chip-multiprocessor, private L1s, shared NUCA L2, based loosely on Compaq Piranha
- `CMP.L2SharedNUCA.OoO` `CMP.L2SharedNUCA.Inorder` with an out-of-order processor core
- `CMP.MT4.L2Shared.Trace` A non-timing chip-multithreaded simulator with shared L2 to study cache miss rates and to construct flex-points for `CMP.MT4.L2SharedNUCA.OoO`
- `CMP.MT4.L2SharedNUCA.OoO` Out-of-order chip-multithreaded processor with private L1s and shared NUCA L2

System Requirements

Operating System

Flexus 4.0.0 runs only under Linux. It has been tested on the following Linux releases:

- SuSE Linux 10.1
- SuSE Linux 8.1
- Fedora Core 8
- Fedora Core 7
- Fedora Core 5

External Software

Flexus 4.0.0 depends on a number of external tools and software libraries:

- GNU GCC 4.1 (or a newer version)
- Boost C++ library 1.33.1
- Virtutech Simics 3.0.X
- GNU make 3.79.1 or newer

GNU GCC 4.x

Flexus uses several advanced features of C++ and requires GCC 4.1 or higher. It will compile with the pre-installed version of GCC on Fedora Core 5, Fedora Core 7 and Fedora Core 8. On other Linux releases, it may be necessary to rebuild GCC from source. Flexus has been tested with GCC 4.1.0 (Fedora Core 5), GCC 4.1.2 (Fedora Core 7 and 8) and GCC 4.1.0 (built from GNU sources).

When building GCC, we suggest using the `--prefix` and `--program-suffix` options of GCC's "configure" script to install GCC 4.1.0 with alternate file names and to an alternate installation directory (rather than install over your system's current gcc installation). We also recommend configuring gcc with `--disablethreads`, as this will improve the performance of Flexus by nearly 30%. Here is the sequence of commands we used to build gcc:

```
(assumes gcc 4.1.0 is untarred in ./gcc-4.1.0)
mkdir gcc-4.1.0-build
cd gcc-4.1.0-build
../gcc-4.1.0/configure --with-prefix=/home/gcc-4.1.0 --with-suffix=-4.1.0
--disable-threads --enable-languages=c,c++
make profiledbootstrap
make install
```

Complete documentation of GCC is available at <http://gcc.gnu.org/>. GCC 4.1 can be downloaded from any of the mirror sites listed at <http://gcc.gnu.org/mirrors.html>. Instructions on installing and building GCC are provided at <http://gcc.gnu.org/install/>. We provide current links to these sites on the SimFlex web page (<http://parsa.epfl.ch/simflex/>). You will need to specify the path and executable name for GCC 4.1 when setting up Flexus.

Boost C++ 1.33.1

The Boost C++ libraries are a set of freely redistributable source code libraries for C++, and are used extensively throughout Flexus. Flexus will not compile with Boost versions other than 1.33.1.

The Boost libraries are documented at <http://www.boost.org/>. We provide a link to the Boost 1.33.1 download site on the SimFlex web page (<http://parsa.epfl.ch/simflex/>).

Once you have downloaded Boost, unpack the distribution to a desired location (e.g., /home/boost). You will need to specify the path where Boost is installed when setting up Flexus. It is not necessary to compile any parts of Boost—the Flexus makefiles will do this for you.

Virtutech Simics

Flexus runs as a loadable module in Simics. We have tested this Flexus 4.0.0 release with Simics 3.0.31. The scripts included with this release will not work with versions of Simics older than Simics 3.x. Versions of Simics 3.x older than 3.0.27 will likely *not* work without modifications to Flexus.

Virtutech provides complete documentation on installing Simics. You should verify that you can run Simics successfully before proceeding with Flexus.

Please contact Virtutech (<http://www.simics.net/>) for more information on obtaining and installing Simics. Personal licenses are available to academic users at no charge (<http://www.virtutech.com/academia/licensing.html>).

Flexus Distribution

The latest Flexus distribution is always available at <http://parsa.epfl.ch/simflex/>.

GLIBC compatibility

Because Flexus is built with a new version of GCC, it requires the latest GLIBC runtime libraries to function properly. Simics is built with an older version of gcc (3.2.2). Hence, Virtutech ships the GLIBC runtime libraries from gcc 3.2.2 with Simics. These older GLIBC libraries are not compatible with Flexus. As part of the Flexus installation process, the Flexus makefiles replace the GLIBC libraries in the Simics distribution with newer libraries. On many Linux distributions (e.g., SuSE Linux 10.1, SuSE Linux 8.1, Fedora Core 8, Fedora Core 7 and Fedora Core 5), both Flexus and Simics will work without difficulty using the newer libraries.

Building Flexus

Preparing the Build Environment

The following procedure will help you setup the Flexus 4.0.0 code and make system to build in your environment.

1. Install GCC, Boost 1.33.1, and Simics 3.0.xx as described above.
2. Unpack the Flexus distribution. Throughout this documentation, we refer to the directory where you unpack the Flexus source as `FLEXUS_ROOT`.
3. Move the `FLEXUS_ROOT/scripts` to your desired location. Throughout this documentation, we refer to the directory where you placed the script directory as `SCRIPT_ROOT`.
4. Modify the path in `SCRIPT_ROOT/global.run_job.rc.tcl`. The scripts for running Flexus need to know the locations where Simics and simulator configurations are located and where the results of the simulation should be placed. You specify these in `global.run_job.rc.tcl`.
 - basedir:** Path where the results of the simulation should be placed. Throughout this documentation, we refer to this directory as `BASE_DIR`.
 - specdir:** Path where the global configuration parameters are located. These configuration parameters are provided in the script distribution. Therefore, it is sufficient that this variable points to `SCRIPT_ROOT/global_job-specs`. This directory contains global configuration parameters for Flexus simulation.
 - userspecdir:** Path where the user's configuration parameters are located (e.g., cache size, associativity, ...). Throughout this documentation, we refer to this directory as `SPEC_DIR`.
 - ckptdir:** Path where the generated flex-points should be placed (will be explained later). Throughout this documentation, we refer to this directory as `CKPT_DIR`.
 - simicsdir:** Path where Simics has been installed.
 - postprocess:** Path to the `postprocess.sh` in the `SCRIPT_ROOT`.
5. Copy `SCRIPT_ROOT/.run_job.rc.tcl` to your home directory. [Note: This file is a hidden file]
6. Modify the paths in `SCRIPT_ROOT/postprocess_ckptgen.sh`.
 - MUNGE:** Path to the `munge-checkpoint.py` file in the `SCRIPT_ROOT`.
 - SIMICS_DIR:** Path where Simics has been installed.
7. Modify the paths in `SCRIPT_ROOT/flexus_scripts.csh` and `SCRIPT_ROOT/flexus_scripts.sh`.
8. Modify the paths in `SCRIPT_ROOT/flexus_scripts.source` to point to the `flexus_script.*` in the `SCRIPT_ROOT`.
9. Modify the paths in `FLEXUS_ROOT/makefile.defs`. The Flexus makefiles need to know the locations where GCC, Boost, and Simics are installed. You specify these in `makefile.defs`. You must set the following four make variables:
 - GCC_PATH** path where GCC 4.0.0 or 4.1.0 is installed. This is the root directory of the GCC installation, which you specified with the `--with-prefix` option when configuring GCC. If using the system's built-in GCC, set this to `/usr`.
 - GCC_BINARY** The name of the `g++` executable. If you build GCC using our recommended options, it will be `g++-4.1.0`. The system's built-in binary is `g++`.
 - BOOST_PATH** Path where the Boost 1.33.1 distribution has been unpacked. This directory contains `boost`, `libs`, and `doc` subdirectories, among others.
 - SIMICS_PATH** Path where Simics has been installed. This directory contains `x86-linux`, `scripts`, and `config` subdirectories, among others.

The Flexus makefiles will verify these paths and check the versions of each external dependency every time the makefiles are used. The version checks can be disabled if you want to use Flexus with unsupported versions of the external tools (see `makefile.checksetup` for details).

4. To install the GLIBC runtime libraries into the Simics distribution, enter the `FLEXUS_ROOT` directory and type:

```
make install
```

Flexus will install GLIBC runtime libraries from GCC 4 into the Simics distribution. See the preceding

discussion of library compatibility for details. This command replaces the GLIBC runtime libraries in `SIMICS_PATH/x86-linux/sys/lib`, and backs the original libraries up in `SIMICS_PATH/x86-linux/sys/lib/`. Running `make uninstall` will cause Flexus to restore the original libraries.

Flexus is now fully installed and ready to build simulators.

Compiling Flexus

Flexus targets are built by issuing the command:

```
make <target>
```

where `<target>` consists of:

```
<component or simulator name>[-<debug level>] [-<subtarget>] [-<architecture>]
```

If you run `make` without any arguments, Flexus will list the components and simulators it knows how to build. The Flexus make system automatically determines the set of simulators and components that can be built by examining the contents of the `components` and `simulators` subdirectories. Thus, any new components or simulators that are added will automatically be recognized by the makefiles. Each simulator includes a makefile called `FLEXUS_ROOT/simulators/<simulator name>/Makefile.<simulator name>` that indicates the set of components required by that simulator. All of a simulators' dependencies are built before the simulator.

The Flexus make command line allows debug statements to be included in the compilation or excluded from it. Each debug statement has an associated severity-level. The optional `-<debug level>` portion of the make target controls the minimum debug severity level included in the build. Debug statements with lower severity are excluded from compilation through preprocessor macros. Debug severity levels are documented in the SimFlex tutorial.

The optional `-<subtarget>` portion of the make target allows special operations or variants of a make target to be built. All make targets support a `-clean` subtarget, which deletes all intermediate object files that may exist for that target, and a `-clobber` subtarget, which additionally removes the output binaries and dependency files created by a target.

In addition to components and simulators, the Flexus make system supports a number of special targets:

- **all** Builds, cleans, or clobbers all simulators and components
- **install** Install compatible GLIBC into Simics
- **uninstall** Restore original GLIBC libraries
- **stat-manager** Build the stat-manager tools

The optional `-<architecture>` portion defines the architecture that Flexus will be targeted. If not defined, Flexus will be compiled for Sparc V9 architectures. X86 targets can be defined by the “-x86” flag.

Here are some example make command lines:

```
make all-clobber
```

Completely removes all intermediate files and binaries from the Flexus source tree. If you experience unexpected “Simics getting shaky errors” when running with Flexus, we suggest clobbering and rebuilding Flexus to ensure that your build is consistent.

```
make CMP.L2SharedNUCA.OoO-iface
```

Builds the `CMP.L2SharedNUCA.OoO` simulator and all its dependencies with `iface` severity debug statements enabled.

```
make CMP.L2SharedNUCA.Inorder-x86-vverb
```

Builds the in-order `CMP.L2SharedNUCA.Inorder` simulator and all its dependencies with `vverb` severity debug statements enabled, for x86 architectures.

Running Experiments in Flexus

The following sections walk through an example timing experiment with Flexus. Over the course of the following sections, you will perform all the steps necessary to prepare a flex-point library and create a report of the results of that experiment, including confidence interval calculations. The example experiment makes use of a variety of scripts included in the Flexus distribution to automate the steps of the process. We suggest using these scripts as a starting point for using your own applications with Simics and Flexus. To complete this experiment, you will require approximately 2GB of free disk space. Any warnings that you might encounter during this process are listed in the Appendix.

The test application is a simple parallelized image processing filter that applies a blur filter to a ppm image. The filter is repeatedly applied to the image, to emulate iterations as typically found in parallel application benchmarks (e.g., SPLASH). Once the simulations are completed, you will estimate the IPC (instructions per cycle) of the application with an associated confidence interval, and construct a breakdown of execution stalls. (Optionally, you may construct an Excel graph depicting the execution time breakdown.)

The complete procedure to prepare and run the sample experiment involves the following steps.

1. Install and compile Flexus.
2. Install the simulated operating system image in Simics
3. Create a Simics checkpoint with the simulated operating system booted and running the `flexus-test-app`
4. Create a flex-point library for detailed simulation of the second iteration of `flexus-test-app`
5. Run out-of-order simulations of the flex-point library
6. Aggregate the statistics databases created by Flexus and generate reports

The subsections below detail each step.

Installing and Building Flexus

Follow the steps in the previous section to obtain the necessary software, install, and build Flexus. To complete the example experiment, you must build the following targets: (see also Appendix:A1)

SPARC simulated architectures

```
make UP.Trace
make UP.OoO
make stat-manager
```

x86 simulated architectures

```
make UP.Trace-x86
make UP.Inorder-x86
make stat-manager
```

OS installation

Before you can run Flexus or any other applications, you must construct a Simics disk image of an operating system.

Installing Solaris

The `flexus-test-app` has been tested with Solaris 10. As of this writing, Solaris 10 CD images can be downloaded from <http://unix.miet.ru/pub/Solaris/iso/sparc/cd/10/>. Simics includes scripts to automate the Solaris installation. Detailed instructions on how to install Solaris under Simics are in Chapter 6 of the Simics/Serengeti Target Guide, pages 56-57. Simics 3.0.30 includes a Serengeti system configuration called “abisko” which simulates an Ultrasparc III system. We have tested `flexus-test-app` using this system configuration. Note that the Solaris installation process may take up to 2 hours.

The installation process creates a Simics disk image called **abisko-sol10-install.disk** and a state file called **abisko-sol10.state**. Verify the checkpoint (last step from the manual on page 57) and place these files in a directory of your choice. Modify the `checkpoint_path` variable in `abisko-sol10.state` to match the path where the disk image resides and change the line `boot_command: "boot -v"` to `boot_command: "boot -rv"`.

Installing Linux on x86 machines

Chapter 6 of Simics x86-440BX Target guide provides instructions on how to install Linux on different x86 architectures. For our example, we downloaded directly from the Simics website a disk dump of a ‘tango’ simulated machine running Fedora Core 5 (<http://www.simics.net/pub/simics/import/x86/tango1-fedora5.craff>). Other Linux disk images are available for download from the Simics website.

Download the disk image into the directory of your choice and verify that you are able to boot the machine with this disk image by running simics with the tango-common.simics script, which can be found in the **SIMICS_ROOT/targets/x86-440bx** directory.

Preparing the Test Application

After generating or acquiring a disk image of the system you want to simulate, you can boot the machine and launch the flexus-test-app. The test application is distributed in source form and as a Simics disk image. The source code and input files are included in **FLEXUS_ROOT/flexus-test-app/src** and **FLEXUS_ROOT/flexus-test-app/inputs**, respectively. You should copy **trace**, **timing_v9**, and **timing_x86** from **FLEXUS_ROOT/flexus-test-app/** to your **SPEC_DIR**. You also need to copy **FLEXUS_ROOT/flexus-test-app/flexus_test_app_v9** and **FLEXUS_ROOT/flexus-test-app/flexus_test_app_x86** folders to your checkpoint directory (**CKPT_DIR**). The test application binary and inputs are preloaded on Simics disk images located in **CKPT_DIR/flexus_test_app_v9/baseline/phase_000/simics** (or **CKPT_DIR/flexus_test_app_x86/baseline/phase_000/simics** for x86 simulated architecture). “**flexus-test-app-disk_v9.craff**” has a compiled Solaris10 version of the application and “**flexus-test-app-disk_x86.craff**” an x86 Linux version.

Before using any Flexus script, you need to initialize environmental variables which are used in all scripts. In **SCRIPT_ROOT**, there is a simple script which does the initialization of the environmental variables. For this purpose, you should run:

```
source SCRIPT_ROOT/flexus_scripts.sourceme
```

You may want to add this command to your shell startup configuration file to be run automatically each time you create a new terminal.

The first step is to create a phase for flex-point generation. Flexus includes a script which automates the process of booting and launching the test application. Go to the **FLEXUS_ROOT/flexus-test-app** directory. Edit the **prepare-test-app** script to select the architecture by setting the corresponding variable **IS_X86_ARCH** or **IS_V9_ARCH** to ‘1’. Edit the **prepare-test-app-xxx.simics** script and set **@ckptdir** to point to your checkpoint directory. To prepare the test application run:

```
./prepare-test-app (see also Appendix:A2, A3)
```

This script launches Simics to boot a single-CPU system and start the test application. If you wish to test a multiprocessor system, modify the **@test_app_threads**, the **\$num_cpus** and the **\$megs_per_cpu** lines in **prepare-test-app-xxx.simics**, and use **CMP.L2Shared.Trace**, **CMP.L2SharedNUCA.OoO**, and **CMP.L2SharedNUCA.Inorder** simulators. Comments in the file indicate the appropriate options for a 4-CPU system.

At the end of the boot process, the scripts will automatically save a checkpoint of the system state in **CKPT_DIR/flexus_test_app_v9/baseline/phase_000/simics** (or **CKPT_DIR/flexus_test_app_x86/baseline/phase_000/simics** for the x86 architecture). The boot process may require 5-10 minutes.

Creating Flex-points

Once you have built the **UP.Trace** simulator and prepared **flexus-test-app**, you can run trace simulation or create flex-points for timing simulation. Each flex-point contains a Simics checkpoint along-side a checkpoint of the Flexus cache hierarchy, branch predictor, and multiprocessor cache coherence state.

The flex-point creation process is carried out by using the **run_job** script. This script commands Flexus to create a flex-points spaced every four million instructions. On a SPARC and x86 uniprocessor systems, the script should create 28 (24 for x86) flex-points. To launch this process, go to **FLEXUS_ROOT**, and for SPARC, run:

```
run_job \  
-postprocess "SCRIPTS_ROOT/postprocess_ckptgen.sh flexpoint 28 test_app" \  
-cfg trace -run flexpoint -local -ckpt-gen UP.Trace flexus_test_app_v9
```

and for x86, run:

```
run_job \  
-postprocess "SCRIPTS_ROOT/postprocess_ckptgen.sh flexpoint 24 test_app" \  
-cfg trace -run flexpoint -local -ckpt-gen UP.Trace flexus_test_app_x86
```

The `-postprocess` option specifies the path and parameters of the script that takes the results of the Flexus simulation and generates the flex-points. In our case, the script name is `postprocess_ckptgen.sh` and its parameters are *flexpoint* which indicates we want to generate flex-points, *28 (24)* which indicates we want to generate 28 (24) flexpoints, and *test_gen* which is just a name for the generated flexpoint set. The `-cfg` option indicates the configuration settings of Flexus which should be loaded from the **SPEC_DIR**. You can define your own configuration by making a new directory with all the required files in it, or you can modify the existing ones. The `-run` option specifies the name of a **rungen** in **.run_job.rc.tcl** file which is located in the home directory to get necessary parameters for the simulation. For example, the distance between flex-points which is four million instructions is specified in the *flexpoint* rungen. The `-local` option indicates the simulation should be run on the local machine (-remote can be used to submit the job to a batch processing system such as Condor or PVS). The `-ckpt-gen` option tells the Flexus that we run this simulation to generate flex-points. UP.Trace is the name of the simulator, and finally `flexus_test_app_v9` or `flexus_test_app_x86` is name of workload located in **CKPT_DIR**.

After 10-15 minutes, all flex-points should be ready. The flex-points are each written to a separate subdirectory under **CKPT_DIR/flexus_test_app/baseline/phase_000/**.

Running a Sampled Timing Simulation

With flex-points created, you are now ready to launch a timing simulation. The `run_job` script, also, automates the process of setting up and launching an individual Flexus simulation. You will use this script to launch simulations for all the flex-points in our sample, as follows (see also Appendix:A5, A6, A7) for SPARC:

```
run_job -ma -cfg timing_v9 -run timing -local -state test_app \  
UP.OoO flexus_test_app_v9
```

or for x86:

```
run_job -cfg timing_x86 -run timing -local -state test_app \  
UP.Inorder flexus_test_app_x86
```

The above command runs separate simulations for flex-points 18 through 28 (or 12 through 24 for x86 simulated machine). This simulation excludes the initialization phase of flexus-test-app from timing measurements only measuring the last 11 generated flex-points. The `-ma` flag is required to launch Simics in micro-architecture mode, which is supported by Flexus only for SPARC simulated machines. The `-state` option specifies the name of a flexpoint that we

want to use for the timing simulation.

The simulations write their output files into a subdirectory under **BASE_DIR**. The subdirectory name is based on the current date and time to allow multiple results to co-exist in the results directory. The `-job name` parameter to the `run_job` script can be used to explicitly set the name of the results subdirectory.

Aggregating and Examining Sampled Results

The timing simulation of each flex-point produces a separate Flexus statistics database. Each statistics database contains a wide variety of statistics (instruction counts, cache event counts, etc.) collected over one or more measurement regions, as controlled by the Flexus configuration. The scripts for the example experiment configure Flexus simulator to collect statistics for regions of 50,000 cycles each. For each flex-point, we will use the first two regions (100,000 cycles) for micro-architecture warmup and report the results collected in the third measurement region.

To report overall results, the individual databases must be aggregated into a single database. The aggregation is accomplished by the **stat-collapse** and **stat-sample** tools. **stat-collapse** selects and combines measurement regions within a single statistics database. **stat-sample** aggregates the output of `stat-collapse` across statistic databases and calculates standard deviations for confidence calculations. Both steps are automated by the `collapse_all.sh` `sample-results` script. To run this script:

```
cd BASE_DIR/results/<name>/flexus_test_app_<v9 or x86>
SCRIPTS_ROOT/collapse_all.sh
SCRIPT_ROOT/sample_and_ipcs.sh `pwd`/flexus_test_app_<v9 or x86>
```

These scripts will create a file called **stats_db.out.gz** in the current directory, which contains the aggregated results from all flex-points. **stat-sample** aggregates results in four ways: sum, average, count, and standard deviation. The results of each of these aggregation methods are included in the output statistics database. (Note: the following examples assume `stat-manager` is available on your path. If not, you must give the full path to **FLEXUS_ROOT**/`stat-manager/stat-manager`. Issue the command:

```
stat-manager list-measurements
```

To view the value of every statistic, summed over the flex-points, issue the command:

```
stat-manager print sum
```

stat-manager is also capable of performing calculations on statistics. For example, to print average IPC for the OoO uniprocessor:

```
stat-manager format-string "<EXPR:{sys-uarch-Commits}/{sys-cycles}>" avg
```

Similarly, for the InOrder uniprocessor:

```
stat-manager format-string "<EXPR:{sys-execute-Commits}/{sys-cycles}>" avg
```

Finally, `stat-manager` can process prepared report files. In these reports, `stat-manager` replaces place-holders for statistics and formulas with their values. To print a report of IPC and its associated 95% confidence interval:

```
stat-manager format FLEXUS_ROOT/flexus-test-
app/reports/ipc_OoO_uniproc.rpt      ".*"      - or -
stat-manager format FLEXUS_ROOT/flexus-test-
app/reports/ipc_InOrder_uniproc.rpt ".*"      - or -
```

To generate a report that includes a breakdown of execution time:

```
stat-manager format FLEXUS_ROOT/flexus-test-app/reports/time_breakdown.rpt  
"sum"
```

The output of the execution time breakdown report can be imported into the Excel spreadsheet in ***FLEXUS_ROOT**/flexus-test-app/reports/time_breakdown.xls* to graph the time breakdown.

To import the data:

1. Open time_breakdown.xls
2. Navigate to the "Data" tab
3. Run the preceding stat-manager command to generate the time breakdown report
4. Copy the output of the report to the Windows clipboard
5. Select cell A1 of the "Data" tab and paste the text of the report
6. (Excel 2002): Click the paste context menu and select "Use Text Import Wizard...". Click finish.
7. (Excel 2003): Select the pasted data, choose "Text to Columns..." from the "Data" menu. Click finish.

The report data should fit the shaded portion of the sheet.

Return to the "Time Breakdown" tab to view the graph. Further stat-manager reports and associated spreadsheets can be based on these examples.

Statistics for x86 simulation

In x86 simulated machines some statistics maintain two versions of counters. The first version has the same name as in SPARC simulated machines, while the second one has an "_x86" suffix. The purpose of this is to count correctly events that appear to be slightly different for x86 architectures, but also maintain the initial version of the counter when this makes sense. For example, "Nodes-execute-Commits_x86" and "Nodes-execute-Commits" count the committed instructions across all processors. An instruction that performs two memory accesses will increase the "Nodes-execute-Commits" counter twice but only once the "Nodes-execute-Commits_x86". Also, table walks made by the hardware will increase the "Nodes-execute-Commits" counter but not the "Nodes-execute-Commits_x86".

Appendix

The following section lists **some** of the warnings and errors that you may encounter during the process of compiling, setting up and running simulations with Flexus and Simics. Most of the warnings described in this section can be easily fixed or safely ignored. Having those as example you can then determine if or how much a different type of warning may affect your simulation results.

A.1 Compiler Warnings

Compiling different components and simulators of Flexus under GCC 4.1.2 you may encounter warnings of the following form, which you may safely ignore.

```
MemoryLoopbackImpl.v9_iface_gcc.o
flexus-4.0.0/components/Common/Transports/MemoryTransport.hpp:65:warning:
`Flexus::SharedTypes::<unnamed>::ExecuteStateTag' defined but not used
flexus-4.0.0/components/Common/Transports/MemoryTransport.hpp: 73:warning:
`Flexus::SharedTypes::<unnamed>::uArchStateTag' defined but not used
flexus-4.0.0/components/Common/Transports/MemoryTransport.hpp:
81:warning: `Flexus::SharedTypes::<unnamed>::MuxTag' defined but not used
flexus-4.0.0/components/Common/Transports/MemoryTransport.hpp:
89:warning: `Flexus::SharedTypes::<unnamed>::BusTag' defined but not used
flexus-4.0.0/components/Common/Transports/MemoryTransport.hpp:
98:warning: `Flexus::SharedTypes::<unnamed>::DirectoryEntryTag' defined
but not used
```

A.2 Warning during boot of the simulated machine

While booting Solaris 10 you may encounter the following warnings, which you can safely ignore.

```
Jun 2 17:13:50 abisko fmd[330]: libpkcs11:
/usr/lib/security/pkcs11_softtoken.so unexpected failure in ELF
signature verification. System may have been tampered with. Cannot
continue parsing /etc/crypto/pkcs11.conf
Jun 2 17:13:50 abisko fmd[330]: libpkcs11:
/usr/lib/security/pkcs11_softtoken.so unexpected failure in ELF
signature verification. System may have been tampered with. Cannot
continue parsing /etc/crypto/pkcs11.conf
```

Also, when booting Fedora Core 5 you may encounter the following warnings, which you can safely ignore.

```
Checking for hardware changes/etc/rc3.d/S05kudzu: line 23: 1143
Segmentation fault /sbin/kudzu $KUDZU_ARGS [FAILED]
```

A.3 Warnings and errors about missing parameters in Flexus

While the Flexus module is loaded in Simics, it initializes itself based on the value of the parameters provided. There are three ways to pass a parameter in Flexus. In order of decreasing priority these are: the configuration file provided in every run, the wiring file in the simulator directory, and finally any default values hard-wired in the code. Some parameters, due to their importance, are required to be set by the user, even if their default value is sufficient. If they are not set, either a warning message will appear on the screen – as the following one – and the simulation will use number in the wiring file (or the one hard-wired in the code), either the simulation will fail with an error message.

```
e.g., <ConfigurationManager.cpp:100> {0}- WARNING: -feeder: CMPwidth
(CMPwidth) was not set in initializeParameters(), from the command
line, or from Simics).
```

Finally, another case where you may see warnings about parameters is when you try to define the value of a non-existing parameter, either because the name has been misspelled or the configuration file used matches a different simulator that uses the specified parameter.

A.4 Warnings from the OoO core

Some infrequent behaviors/instructions of the Out-of-Order Ultrasparc core have not been implemented and are either simulated partially or not at all. This doesn't affect the correctness of the simulation, because after committing each instruction the Flexus and Simics states are compared and the Simics state is copied into Flexus if a difference is found. The synchronization process outputs warnings to indicate this event.

A.5 Warnings while trying to stop simulation

Flexus automatically tries to stop simulation and save results when it has reached the target number of simulated cycles. However, at that specific point Simics may not be able to be halted, requiring some more cycles (less than 1K cycles) before it can safely stop. This causes Flexus to issue the exit message multiple times, including multiple messages about successfully saving the results. The file keeping the results will not be affected by this behavior.