

SPEC – Power and Performance

User Guide

SPECpower_ssj2008 V1.10

Standard Performance Evaluation Corporation

Table of Contents

1	Introduction	4
1.1	Preface	4
1.2	General Concepts	4
1.3	The SPECpower_ssj2008 Benchmark Suite	4
1.4	The Hardware Components	5
2	Installation and Setup of SPECpower_ssj2008	6
2.1	Hardware Setup	6
	Power Analyzer	6
2.2	Software Installation	7
	Java Runtime Environment (JRE) Set-up	7
2.3	Network Setup	9
	System Under Test (SUT)	9
	System Under Test (SUT)	9
2.4	Running the JVM Director Remotely	10
	The JVM Director Run Script	10
	Remote JVM Director: SUT	10
	Remote JVM Director: Controller System	10
	Controller System	10
2.5	Trial Run of SPECpower_ssj2008	11
	Controller System	11
	Start the Trial Run	12
	System Under Test (SUT)	12
3	Running SPECpower_ssj2008	17
	Running SPECpower in a multi-node configuration	17
	Inventory of Operating System Services	18
	Results	18
4	Operational Validity	19
5	Metric	19
5.1	SPECpower_ssj2008 Metric	19
6	Results Reports	19
6.1	Overview	19
6.2	The CSV file	20
	CSV File Naming Convention	20
	CSV File Format	20
6.3	The RAW File	22
	Raw File Format	22
6.4	The Reporter	23
	Running the Reporter Manually	23
6.5	The HTML File(s)	24
	HTML File Naming Convention	24
	Overall HTML Results File Format	24
6.6	The SSJ Results Files	25
	The RAW Files	25
	The LOG File	26
	The RESULTS File	26
	The TXT File	26

7	Visual Activity Monitor (VAM)	26
	Global Settings	29
	Screen Size and Location	29
	Title Panel and Text.....	29
	Chart Panel Example	30
	Freeze Button Example.....	30
	Colors 30	
8	Performance Tuning	32
9	Submitting Results	33

1 Introduction

1.1 Preface

SPECpower_ssj2008 is a benchmark product developed by the Standard Performance Evaluation Corporation (SPEC), a non-profit group of computer vendors, system integrators, universities, research organizations, publishers, and consultants. It is designed to provide a view of a server system's power consumption running Java server applications.

This practical guide is intended for people that wish to setup and run the SPECpower_ssj2008 benchmark in order to accurately measure the power consumption of their server in relation to the server's performance.

Updates and the latest version of this User Guide can be found here: http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-User_Guide.pdf.

In order to submit SPECpower_ssj2008 results, the licensee must adhere to the rules contained in the Run and Reporting Rules (http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Run_Reporting_Rules.html) which is included in the benchmark kit.

Be sure to read, understand and follow all the safety rules that come with your Power Analyzer and compute equipment

1.2 General Concepts

1.3 The SPECpower_ssj2008 Benchmark Suite

SPECpower_ssj2008 consists of three main software components:

- Server Side Java (SSJ) – Workload
 - The SSJ-Workload is a Java program designed to exercise the CPU(s), caches, memory, the scalability of shared memory processors, JVM (Java Virtual Machine) implementations, JIT (Just In Time) compilers, garbage collection, and other aspects of the operating system of the SUT.
 - For more information on [SSJ](#) -> http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Design_ssj.pdf
- Power and Temperature Daemon (PTDaemon)
 - The Power and Temperature Daemon is to offload the work of controlling a power analyzer or temperature sensor during measurement intervals to a system other than the SUT.
 - For more information on [PTDaemon](#) -> http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Design_ptd.pdf
- Control and Collect System (CCS), including the Visual Activity Monitor (VAM)
 - CCS is a multi-threaded Java application that controls and enables the coordinated collection of data from multiple data sources such as a workload running on a separate SUT (System Under Test), a power analyzer, and a temperature sensor.
 - VAM is a software package designed to display activity from one, two, or three SUT's simultaneously, in conjunction with the SPECpower_ssj2008 benchmark.

- o For more information on CCS and VAM ->
http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Design_ccs.pdf

1.4 The Hardware Components

The following components work together to collect server power consumption and performance data by exercising the SUT with a predefined workload.

The most basic SPECpower_ssj2008 test bed implementation consists of these four hardware devices:

- Server Under Test (SUT)
The SUT is the system that will be driven by the SSJ workload. The SUT's performance and power consumption characteristics will be captured and measured by the benchmark.
- Power Analyzer
The power analyzer is used to measure and record the power consumed by the SUT.
- Temperature Sensor
The temperature sensor is used to capture the temperature of the environment where the SUT is being benchmarked.
- Controller System
The controller system is a separate system that runs the CCS and Power and Temperature Daemon (PTD) portions of the SPECpower_ssj2008 benchmark suite.

Figure 1.2-1 illustrates the architecture of a simple SPECpower_ssj2008 benchmark implementation.

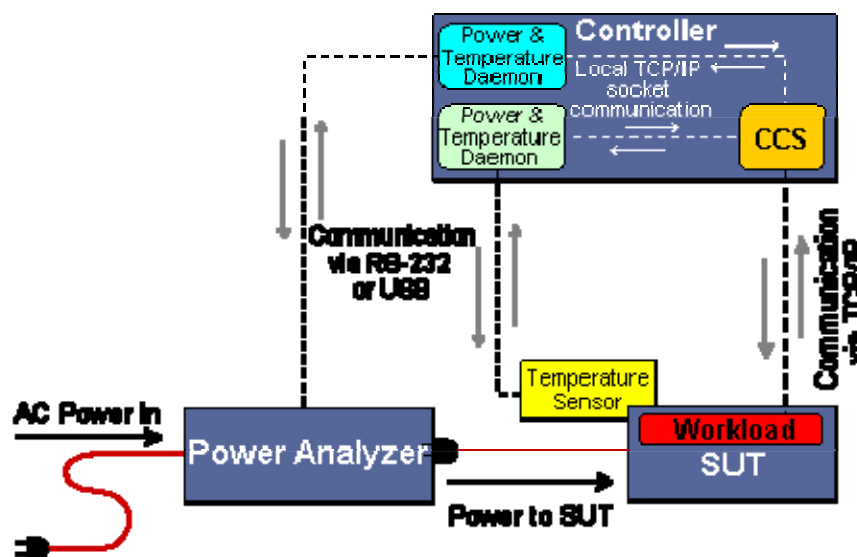


Figure 1.2-1

2 Installation and Setup of SPECpower_ssj2008

This section will go over the installation and setup procedures for a basic SPECpower_ssj2008 implementation consisting of:

- one server under test
- one power analyzer and one temperature sensor
 - The PTD can connect to and collect data from either a power analyzer or a temperature sensor, but never more than one device at a time per PTD instance.
- one benchmark controller system

2.1 Hardware Setup

2.1.1.1 Power & Temperature Daemon (PTDaemon)

In this simple SPECpower_ssj2008 implementation, it is assumed that the PTD instance for the power analyzer and the PTD instance for the temperature sensor are residing locally on the controller server as illustrated in figure 1.3.3. By default, the run scripts for PTD will run the PTD executable in "dummy mode". Dummy Mode is a mode of operation for PTD in which it will return false power and temperature readings during the benchmark run. Like the SSJ and CCS modules of the SPECpower_ssj2008 benchmark, the PTD module also uses run scripts. Therefore, two concurrent instances of PTD must be run; one for communication with the power analyzer, and one for communication with the temperature sensor. These run scripts are recommended for starting the PTD modules.

There is a run script for initializing an instance of PTD for a power analyzer named `runpower.bat` for Windows or `runpower.sh` for Unix. There is a run script for a temperature sensor named `runtemp.bat` for Windows or `runtemp.sh` for Unix. It should be noted that PTD also supports General Purpose Interface Bus (GPIB) connectivity.

Modifications of these scripts are done the same way they are done for the SSJ and CCS run scripts. The variables that the user may be interested in modifying are as follows:

- *The model of power analyzer or temperature sensor to be used in the test bed implementation.*

The device model, for the temperature sensor or for the power analyzer, is specified with an integer value.

To find out which integer value corresponds to the device being used within the test bed, Use the command prompt to change to the directory of the PTD executable, and run the PTD executable for your operating system without arguments.

This will cause PTD to display a list of supported power and temperature devices, along with their corresponding integer values, and find and make note of the values that apply to the devices being used within the test bed

This will cause PTD to display a list of supported power and temperature devices, along with their corresponding integer values. Find and make note of the values that apply to the devices being used within the test bed.

Power Analyzer

Before selecting a power analyzer for the SPECpower_ssj2008 test bed, please consult the *Power Analyzer Minimum Requirements*.

Please consult the manufacturer's documentation for instructions that are specific to the exact analyzer used.

1. Once you have correctly associated the Device number with the corresponding integer values that apply to those devices
2. Connect the communication interface of the power analyzer to the appropriate port of the controller server with a manufacturer-approved cable.
3. Open `runpower.bat` for Windows or `runpower.sh` for Unix with a text editor, and locate the `set DEVICE=` for the `runpower` script respectfully. Enter the correct device number for the power analyzer being used.
 - o `set DEVICE=8` (for Windows & Unix)
4. Now enter the `DEVICE_PORT` that the PTD uses to communicate with Controller system.
 - o `set DEVICE_PORT=COM1` (for Windows)
 - o `set DEVICE_PORT=/dev/tty2` (for Unix)
5. To verify what port the power meter is using you may left-click my computer, click Manage. Click Device Manager and then click Ports this should give you a list of all the ports attached to your computer at that time. Verify the port the power analyzer is connected to.
6. Set power analyzer communication settings as required by the Hardware Setup Guide (http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Hardware_Setup_Guide.pdf).
7. Connect the SUT's AC power cord to the power receptacle provided by the power analyzer.

2.1.2 Temperature Sensor

Before selecting a temperature sensor for the SPECpower_ssj2008 test bed, please consult the *Temperature Sensor Minimum Requirements*.

The temperature sensing device must be secured 50mm from the center of the inlet for airflow of the equipment being benchmarked, please refer to *SPECpower_ssj2008 Run and Reporting Rules*.

Please consult the manufacturer's documentation for instructions that are specific to the exact sensor used.

1. Place the temperature sensor in accordance with the rules outlined in this section.
2. Connect the temperature sensor's communication interface to the appropriate port on the controller server with a manufacturer-approved cable.
3. Open `runtemp.bat` for Windows or `runtemp.sh` for Unix with a text editor, and locate the `set DEVICE=` for the `runtemp` script respectfully. Enter the correct device number for the temperature sensor being used.
 - o `set DEVICE=1001` (for Windows & Unix)
4. Also locate the `set DEVICE_PORT=` for the `runtemp` script respectfully. Enter the correct port number used to connect the temperature sensor.
 - o `set DEVICE_PORT=USB` (for Windows)
 - o `set DEVICE_PORT=/dev/tty1` (for Unix)
5. To verify what port the temperature sensor is using you may left-click my computer, click Manage. Next click Device Manager and then click Ports this should give you a list of all the ports attached to your computer at that time. Find the port that the temperature sensor is connected to.

2.2 Software Installation

Java Runtime Environment (JRE) Set-up

Before proceeding it is the responsibility of the user to select a suitable Java Runtime Environment (JRE) (version 1.5 or later) is installed on the SUT. There are a variety of JREs available from a number of different vendors.

Also note: The JRE is the only auxiliary software that must be installed within the benchmark environment with the SPECpower_ssj2008 benchmark suite. No additional database or web server components are required.

The installation program is written in Java, and will therefore require a Java interpreter in order to run. To invoke the setup program:

1. Change to the root directory of the installation CD.
2. Call the command:
`java -jar setup.jar [options]`

The options to pass into the setup program are as follows:

To run in GUI mode (for Windows), simply type:

```
java -jar setup.jar
```

To run in CONSOLE mode, type:

```
java -jar setup.jar -i console
```

To run in CONSOLE + SILENT mode (where the installer option panels are bypassed and the benchmark will be installed in the default install dir, /SPECpower_ssj2008 on Unix or <SYSTEM_DRIVE_ROOT>:\SPECpower_ssj2008 on Windows):

```
java -jar setup.jar -i silent
```

To run in CONSOLE + SILENT mode with an install dir of the user's choosing, for example, /power_ssj2008:

```
java -jar setup.jar -i silent -DUSER_INSTALL_DIR=/power_ssj2008
```

The exact location within the OS directory structure into which the SPECpower_ssj2008 suite is to be installed is subject to the user's discretion. The user must either add the path to the Java executable to the PATH environment variable (consult the documentation of the OS for instructions), or invoke the Java interpreter with the fully qualified path and filename.

For example:

```
C:\Java\AcmeSuperJava-v0.15\bin\java.exe -jar setup.jar
```

or

```
/usr/java/AcmeSuperJava-v0.15/bin/java -jar setup.jar -i console
```

as appropriate to the OS being used.

To define the environmental variables: Right-click My Computer, select Properties, Select the Advanced tab, select Environmental Variables, select PATH, and click edit. At the end of the list, place the path to the `java.exe` of the `bin` directory that the Java software was installed. IE: `C:\Java\AcmeSuperJava-v0.15\bin\java.exe`. Also add `.JAR` to the `pathext`.

An easy way to test for this is simply to type `java -fullversion` from a command prompt (and from within a directory that does NOT contain the Java executable). If the Java executable is indeed specified in the PATH environment variable, then the output of this command should specify the version of the JRE currently installed on the SUT. If it instead returns an error, then the path to the Java executable is not correctly specified in the PATH environment variable.

2.3 Network Setup

System Under Test (SUT)

For a valid SPECpower_ssj2008 benchmark implementation, there is a set of physical environment criteria that must be satisfied. Please consult the Run and Reporting Rules in order to ensure that the environment in which the benchmark is taking place meets the minimum requirements.

Ensure that power is properly routed to the SUT through the power analyzer *only*.

Please note: the SUT cannot receive power from any source other than that which is provided by the power analyzer. For example, the other power supply/supplies may not be connected to any power source other than that which is provided by the power analyzer. Refer to Quick Start guide

Connect the network interface of the SUT so that it will establish a TCP/IP connection with the controller server.

Configure the operating system's TCP/IP implementation such that the SUT is on the same subnet as the controller. Also disable the firewall for both systems to allow them to communicate through TCP/IP. Try Ping to see if the IP of the SUT to see if they can converse through TCP/IP. Consult the operating system's documentation for specific instructions regarding this procedure. In consideration of the data sources that CCS will be connecting to, ensure that the controller server's TCP/IP interface(s) is configured so that reliable TCP/IP communication can be established with all data sources.

System Under Test (SUT)

The SPECpower_ssj2008 software uses the Java interpreter to run, so the user must ensure that the java interpreter can be found by the benchmark. There are two ways to accomplish this:

Add the fully qualified path of the Java executable to the run script. (Note that the user may also elect this method for the purpose of overriding the Java executable path that is specified in the PATH environment variable.)

1. Change to the SSJ directory and open runssj.bat (for Windows) or runssj.sh (for Unix) with a text editor.
2. For runssj.bat, locate the line set JAVA=java and replace java with the fully qualified path to the java executable. For example:
`set JAVA=c:\Java\AcmeSuperJava-v5.30\bin\java.exe`
3. For runssj.sh, locate the line JAVA=java, and replace java with the fully qualified path to the java executable. For example:
`JAVA=/usr/java/AcmeSuperJava-v5.30/bin/java`
4. Save and close the file.

The recommended (and simplest) method for starting the SSJ iteration is to use the sample run script that came with the SPECpower_ssj2008 distribution.

1. Change to the directory in which the SSJ workload module was installed.
2. Run the file, `runssj.bat` or `runssj.sh`, as appropriate to your operating system.

Please note: it is the responsibility of the user to inspect the sample run script files for any commands, parameters, or variables that may need to be modified for the particular environment under which the workload is being run.

For Windows, when the workload instance is ready, it will display the text `SSJ instances ready for connection`. This means that the SSJ workload is now waiting for the command from CCS to start.

For Unix, the `STDOUT` is not output to the console by default. Nonetheless, running `runssj.sh` will start the SSJ iteration.

2.4 Running the JVM Director Remotely

As mentioned, the JVM Director need not necessarily be run on the SUT. This section will very briefly go over the steps to configure the test bed for running the JVM Director on the controller system. This method is an alternative from the default method, and is of course entirely optional.

The JVM Director Run Script

Like most other components of SPECpower_ssj2008, a run script is provided with the kit with which to initiate the JVM Director. The default name of the script is `rundirector.bat` or `rundirector.sh`, and it is located under `.../SPECpower_ssj2008/ssj`. To start the director, simply call the `rundirector` script from the command line.

In order for Director to successfully connect to the SUT, the SUT must know the IP of the Director:

Remote JVM Director: SUT

1. On the SUT, open the SSJ run script (`runssj.bat` or `runssj.sh`) for editing.
2. Locate the `LOCAL_DIRECTOR=` property. It should equal `FALSE`, if the Director is on another system change `FALSE` to `TRUE`, and fill in the IP address of the `DIRECTOR_HOST` to the controller IP address used.
3. Save and close the SSJ run script.
4. Open the control properties file, `SPECpower_ssj_EXPERT.props`, for editing.
5. Uncomment (remove '#') the property `input.director.hostname`, and change its value from `localhost` to the hostname or IP address of the system that will be running the JVM Director (for this example, the controller system).
6. Save and close `SPECpower_ssj_EXPERT.props`.
7. Open the descriptive properties file, `SPECpower_ssj_config.props`, for editing.
8. Change the property `config.director.location` to reflect the system on which the JVM Director will be running. In this example, `Controller`.
9. Save and close `SPECpower_ssj_config.props`.

Remote JVM Director: Controller System

1. Open the CCS properties file, `ccs.props`, for editing.
2. Change the property `ccs.wkld.ssj_dir.IP` to the hostname or IP address of the system that will be hosting the JVM Director. For this example, the controller system which is `localhost` or `127.0.0.1`.
3. Save and close `ccs.props`.

Controller System

Please see *SPECpower_ss2008j Run and Reporting Rules* for a list of minimum system requirements for the controller server. Connect the network interface of the controller server so that it can establish a TCP/IP connection with the SUT, refer to section 2.4. System Under Test (SUT).

Before proceeding it is the responsibility of the user to select a suitable Java Runtime Environment (JRE) (version 1.5 or later) is installed on the Controller system. There are a variety of JREs available from a number of different vendors. It is the responsibility of the user to select a suitable JRE for the benchmark environment. (Please refer to section 2.2.1 Java Run-time Environment)

2.5 Trial Run of SPECpower_ssj2008

After successful installation of the SPECpower_ssj2008 benchmark suite on the Controller System and the SUT, the user may be interested in starting a trial run of the benchmark. This will ensure that all of the fundamental components of the test bed are functioning properly. This section will go over the steps involved in doing so. Please note that the CCS and SSJ modules will, by default, wait for 300 seconds (5 minutes) for a connection with the JVM Director. If the JVM Director is for some reason not initiated sooner than 300 seconds after either the SSJ workload module or the CCS module is initiated, then the module will time-out and terminate with an error and will then need to be restarted.

2.5.1.1 Director

As was done for the SSJ modules on the SUT, the user must ensure that the Java executable is accessible for the Director module on the Controller System as well. The Java interpreter must be accessible via either the `PATH` environment variable of the OS, or by specifying the fully qualified path to the Java executable in the run script.

1. Change to the SSJ directory within the benchmark suite and open `rundirector.bat` (for Windows) or `rundirector.sh` (for Unix) with a text editor.
2. For `rundirector.bat`, locate the line `set JAVA=java`, and replace `java` with the fully qualified path to the Java executable. For example:
`set JAVA=c:\Java\AcmeSuperJava-v5.30\bin\java.exe`
3. For `rundirector.sh`, locate the line `JAVA=java`, and replace `java` with the fully qualified path to the Java executable. For example:
`JAVA=/usr/java/AcmeSuperJava-v5.30/bin/java`
4. Save and close the file.

2.5.1.2 Control and Collect System (CCS)

Controller System

The controller server runs the CCS module of the SPECpower_ssj2008 benchmark suite. CCS is a Java based application, therefore the user should ensure that a suitable Java Runtime Environment (version 1.5 or later) is properly installed on the controller server before proceeding.

The user must ensure that the Java executable is accessible for the CCS module on the Controller System as well.

1. Change to the CCS directory within the benchmark suite and open `runCCS.bat` (for Windows) or `runCCS.sh` (for Unix) with a text editor.
2. For `runCCS.bat`, locate the line `set JAVA=java`, and replace `java` with the fully qualified path to the Java executable. For example:
`set JAVA=c:\Java\AcmeSuperJava-v5.30\bin\java.exe`
3. For `runCCS.sh`, locate the line `JAVA=java`, and replace `java` with the fully qualified path to the Java executable. For example:
`JAVA=/usr/java/AcmeSuperJava-v5.30/bin/java`

4. Save and close the file.

Start the Trial Run

Now the trial run is ready to be started:

1. Use the run script to start the PTD instance for the temperature sensor by changing to the PTD directory and running `runtemp.bat` or `runtemp.sh` as appropriate to the OS. This will start the Temperature daemon instance.
2. Use the run script to start the PTD instance for the power analyzer by changing to the PTD directory and running `runpower.bat` or `runpower.sh` as appropriate to the OS. This will start the Power daemon instance.
3. On the SUT, change to the SSJ directory, and run `runssj.bat` (for Windows) or `runssj.sh` (for Unix). SPECpower_ssj2008 workload iteration will start up, and wait for a connection from the Director.
4. On the controller server, change to the SSJ directory, and run `rundirector.bat` (for Windows) or `rundirector.sh` (for Unix).
5. Next on the controller server, change to the CCS directory, and use the run script `runCCS.bat` or `runCCS.sh`, as appropriate to the OS, to start the CCS iteration.

This will start the SPECpower_ssj2008 trial run, and will take approximately 73 minutes to complete. The purpose of the trial run is to identify possible issues with the benchmark test bed implementation. If the procedures in this section were followed correctly, and problems are still encountered during the trial run, then issues are native to the test bed that must be resolved in order for SPECpower_ssj2008 to run properly.

It should be understood that the number of possible complications that can be encountered with any test bed implementation can be quite numerous, ranging from OS installation issues, to network problems, to Java Runtime Environment issues. The entire spectrum of troubleshooting procedures for any server/networking environment, are immense therefore troubleshooting procedures for any particular test bed implementation fall outside the scope of this document.

Software Setup

Having done a trial run of the benchmark, the user should now learn what the changeable parameters are, and where to find and change them in the run scripts and properties files.

System Under Test (SUT)

2.5.1.3 Properties Files: Overview

The SSJ module takes two properties files as input:

- A control properties file: The control properties file is used to modify the operation of the benchmark descriptive properties file
 - o The descriptive properties file is used to document the SUT, documentation is required for publishable results, and is reflected in the output files.

The control properties files come in two different forms: basic and advanced. It is up to the user to decide which one to use for the benchmark run (this will be covered in

greater detail in the next section). The values of the descriptive properties file do not affect the workload. The default file name for the descriptive properties file is `SPECpower_ssj_config_sut.props`.

Before modifying these files, it is good practice to first make a copy of the originals to enable easy recovery of the default benchmark configuration. The control properties file can be renamed to any name desired by the user providing that the user updates the `set PROPROFILE=`, or the `PROPROFILE=`, variable in the `runssj.bat` or `runssj.sh` file. The descriptive properties files can also be changed to any name desired by the user. Exactly how this is done is explained in the next section.

2.5.1.4 Properties Files: Format

The properties files are in plain-text ASCII format, and can be easily edited with most popular text editors. Each line of a properties file is one of three types:

- a blank
- a comment
Any text following a # character will be ignored by the benchmark
- a property assignment
Each property assignment is of the form `name=value`, where "name" is the property identifier (the terms "name" and "property" are used interchangeably).

Property names are specific to the benchmark and *must not be changed*. "value" is the actual value that is assigned to the property. A value can take the form of an integer, a string of text, or a boolean value depending on the property type. The SSJ workload engine takes these property values as input for the benchmark iteration.

As mentioned previously, the control properties file comes in two forms:

- a basic control properties file whose default name is `SPECpower_ssj.props`
- an advanced control properties file whose default name is `SPECpower_ssj_EXPERT.props`

The basic control properties file contains just a few basic parameters for the SSJ run. This properties file is more suitable for beginners, or simply for benchmark implementations that are relatively simple.

The advanced properties file contains quite a large number of additional parameters that the user may elect to modify for the benchmark run.

Please note that there are also a number of control properties whose values *must not be changed* in order for the SPECpower_ssj2008 result to be publishable. See the *Run and Reporting Rules* for more information.

2.5.1.5 Control Properties: SSJ

One of these files must be selected for use with the benchmark run by setting the `PROPROFILE=` variable in the SSJ run.

For simplicity, only those parameters that are the most pertinent to the discussion of the basic concepts of the SSJ workload will be covered in this section.

The properties are contained in basic `SPECpower_ssj.props` file contains only the following:

- `input.load_level.number_warehouses`
 - `input.calibration.interval_count`
 - `input.director.hostname`
 - `input.status.port`
 - `input.include_file`
 - `input.output_directory`
-
- *Length of the run*

The user can change the amount of elapsed time that it will take to complete the run. It should be understood that while these parameters can be adjusted, there are restrictions on which values can be altered, and how much they can be altered for publishable results. Please see *SPECpower_ssj2008 Run and Reporting Rules* for details.

 - `input.load_level.length_seconds`: This is the number seconds that each target load will run for. The default is 240. Remember, the benchmark run will iterate through several target loads; thus for m target loads and n seconds, $m*n$ total seconds will be added to the length of the run.
 - `input.calibration.interval_count`: This specifies the number of calibration intervals that the benchmark will iterate through before the measurement intervals. More calibration intervals will increase the run time of the benchmark.
 - `input.calibration.length_seconds`: This is the number of seconds that each calibration interval will run for. The default is 240 seconds.
 - `input.idle.length_seconds`: This is number of seconds that the active idle measurement interval will run for. The default is 240 seconds.
 - *How the workload instance(s) communicate with the JVM Director*

The JVM workload instance(s) communicate with the JVM Director via TCP/IP socket communication. The JVM Director can be run locally on the SUT, or it can be run remotely.

 - `input.director.hostname`: Use this property to inform the SSJ workload engine of the network hostname of the system running the JVM Director instance. By default, the JVM Director will run locally on the SUT, and the default value of this property will be "localhost". This value can take the form of an IP address as well.
 - *How the benchmark reports results*

`input.include_file`: Use this property to specify the name of the descriptive properties file that is to be used when writing the results reports. The default is `SPECpower_ssj_config.props`.

 - `input.output_directory`: Use this property to specify the directory into which SSJ will store the results reports, and the user must ensure that the SSJ module has sufficient access and write credentials for the directory specified.
 - *Number of concurrent threads / warehouses*

There is a one to one relationship of threads to warehouses.

 - `ccs.load_level.number_warehouses`: Modify this property in order to change the number of concurrent warehouses (and consequently the number of concurrent threads) that will be run per JVM instance throughout the benchmark iteration.
 - *How the JVM Director communicates with CCS*

The JVM Director communicates with CCS on the controller server via TCP/IP socket communication.

- o `input.director.enabled`: To enable or disable communication with CCS, set this property to true or false. If set to false, the workload will run in stand-alone mode (without intervention from CCS. Used only for testing or debugging purposes). If set to true (necessary for a valid run), then the workload iteration will not start until instructed to do so by CCS, and will subsequently continue to be polled and controlled by CCS throughout the benchmark iteration.
- o `input.director.port`: Use this property to specify the TCP port on which CCS will attempt initiate communication. This will tell the JVM Director what TCP port to listen on for CCS.

2.5.1.6 Descriptive Properties

The descriptive properties file contains all of the information about the SUT including the hardware configuration, software and operating system configuration, benchmark configuration, tuning information and any other notes about the SUT. The properties in this file must be accurate for a valid, publishable run. As mentioned previously, the default filename of the descriptive properties file is `SPECpower_ssj_config.props`. This file is called by the *reporter* upon completion of the benchmark iteration, and is used by the reporter to generate the results output files. The reporter and output files will be explained in further detail in section 6. The descriptive properties file is extremely self-explanatory. More information about the descriptive properties file can be found in appendix A.

2.5.1.7 Network Configuration: System Under Test (SUT)

Refer to section 2.4 to configure the operating system's TCP/IP implementation such that the SUT can communicate with the controller server. Consult the operating system's documentation for specific instructions regarding this procedure.

2.5.1.8 Control Properties: Control and Collect System

Similar to the architecture of the SSJ module on the SUT, the CCS module reads a plain text properties file for its control configuration. In the case of CCS, however, no descriptive properties file is used. Instead, the descriptive properties that are particular to the CCS and PTD implementation are also contained in the control properties file. The default name for the properties file is `ccs.props`. Just like SSJ's properties files, the CCS properties file may be renamed at the user's discretion, provided the proper variable is updated in the run script. The `ccs.props` file is used to control all configurable parameters of the CCS module of SPECpower_ssj2008, including:

- *The names of the data sources to which CCS will connect*
There are three types of data sources that CCS can connect to and collect data from: An SSJ workload instance (on the SUT), a power analyzer (via a PTD instance) and a temperature sensor (via an additional PTD instance).

`ccs.ptd`: Set this property to define the name of the power analyzer and temperature sensor data sources. Refer to SPEC_power-design_CCS section 2.1.2. The syntax for specifying these data sources is `source1, source2`. For example:

```
ccs.ptd = pwr1, temp1
```

- *TCP/IP communication with the data sources*
CCS utilizes TCP/IP socket connections in order to communicate with the other SPECpower_ssj2008 modules within the test bed. In order for CCS to establish these communication paths, the TCP/IP information for each data source must be specified. For example: Suppose the SSJ workload data source was assigned the name MyWorkload with the property assignment
`ccs.wkld = MyWorkload`

Now that this workload has been defined, CCS needs to know how to establish a connection with the workload. Assume that the IP address of the SUT is 192.168.1.3 and that the JVM Director on the SUT has been configured to listen for CCS on TCP port 8886 (default). To achieve communication, enter the following properties like this:

```
ccs.wkld.MyWorkload.type = SSJ
ccs.wkld.MyWorkload.IP = 192.168.1.3 or localhost
ccs.wkld.MyWorkload.Port = 8886
```

Where the JVM Director is to be run is subject to the sole discretion of the user.

For this example test bed, define the other two data sources: the power analyzer and the temperature sensor.

Use the `ccs.ptd` property to define the power analyzer and temperature sensor data sources respectively:

```
ccs.ptd = MyPower,MyTemp
```

Now that these data sources have been defined, CCS needs to know how to connect to the data sources. This is achieved in the same manner in which the SSJ connection information was defined:

```
ccs.pwr.MyPower.type = PowerMeter
ccs.pwr.MyPower.IP = localhost
ccs.pwr.MyPower.Port = 8888

ccs.temp.MyTemp.type = TempSensor
ccs.temp.MyTemp.IP = 127.0.0.1
ccs.temp.MyTemp.Port = 8889
```

Since, for this example, both PTD instances for the power analyzer and the temperature sensor are being run locally on the controller server (remember, they may be run remotely if so desired), both network paths are set to the local machine (`localhost` or `127.0.0.1`). Both the hostname, `localhost`, and the IP address, `127.0.0.1`, are used in this example to illustrate that either a host name or an IP address may be used in the properties file.

The sample `ccs.props` file that comes with the distribution already contains sample properties for the different data sources. The user need only modify the existing sample properties for a benchmark setup.

- *The descriptive properties of the CCS and PTD implementation.*
The descriptive properties contained in the `ccs.props` file are very self-explanatory. Below is an excerpt from a sample `ccs.props` file:

```
ccs.config.hw.vendor=Hoover
ccs.config.hw.model=SuperStation Deluxe
```

```
ccs.config.hw.cpu=StrongARM SA-110
ccs.config.hw.cpu.characteristics=200MHz 16KB/16KB MMU
ccs.config.hw.memory.gb=2
ccs.config.sw.os=Amiga OS
ccs.config.sw.jvm.vendor=Acme
ccs.config.sw.jvm.version=AcmeSuperJava 5.30

ptd.pwr1.config.analyzer.vendor=Eureka
ptd.pwr1.config.analyzer.model=Power Cyclone
ptd.pwr1.config.analyzer.serial=1234567
ptd.pwr1.config.analyzer.connectivity=Serial Port
ptd.pwr1.config.calibration.institute=NIST
ptd.pwr1.config.calibration.accredited_by=Jane
ptd.pwr1.config.calibration.label=NIST label
ptd.pwr1.config.calibration.date=Jan-1-07
ptd.pwr1.config.ptd.system=Controller System
ptd.pwr1.config.ptd.os=Amiga OS

ptd.temp1.config.sensor.vendor=Acura
ptd.temp1.config.sensor.model=NS Thermal X
ptd.temp1.config.sensor.driver=1.2.3.4
ptd.temp1.config.sensor.connectivity=USB
ptd.temp1.config.ptd.system=Controller System
ptd.temp1.config.ptd.os=Amiga OS
```

Modify each descriptive property so as to accurately reflect the CCS and PTD implementation.

3 Running SPECpower_ssj2008

After completing the setup procedures in section 2, the benchmark is ready to run.

If the user has not already done so, it is a good idea to initiate a trial run of the benchmark to ensure that SPECpower_ssj2008 benchmark suite is installed and configured correctly. See [section 2.5](#) for instructions on how to initiate a trial run.

Running SPEC_power in a multi-node configuration

1. Find two power analyzer or temperature sensors certified by SPEC. Please consult the manufacturer's documentation for instructions that are specific to the exact analyzer used.
2. Connect the power analyzer or temperature sensor using RS-232 or USB connections. Make sure the hub that you use is OS supported. This would make it easier to connect the two power analyzer or temperature sensors to your controller system.
3. Change to your CCS directory and find `CCS.props`. Open `CCS.props` in a text editor and locate Data Source Entries, under this section there is a field named `CCS.ptd`. This field allows CCS to know the devices that are trying to communicate with the controller system and which ports the devices are attempting to communicate on. In this field add to it the same number of power analyzers that you are directly connecting to the controller. Ex: If the power meter is the second power analyzer then in that field you will add `pwr2`, and it would be the same for the temperature sensor you would add `temp2`.
4. Notice the detail section that is located in the `CCS.props` file. In this section will be the field of the device that you added to the `CCS.ptd` variable. Uncomment the `CCS.ptd.device.type`, `CCS.ptd.device.IP`, and `CCS.ptd.pwr2.port` this

- informs the CCS how to communicate with the additional device. Save `CCS.props` and exit.
5. To connect to the additional power or temperature sensor you can make a copy of the `runpower` or `runtemp` script. Open the script up in a text editor to edit the script. Verify the field `set NETWORK_PORT` corresponds to the port that is assigned in the device in `CCS.props`. Alter the device port used to communicate with the device and if a different power analyzer is used, please change the device type.
 6. Now run the `runpower` or `runtemp` scripts so that you make sure your connections are correct. If they are correct then you will receive a message from both scripts stating "Waiting for connections". If not, then check the run scripts for any discrepancies dealing with the network or device ports.
 7. In order for the controller and the SUT to communicate, both must be on the same subnet.
 8. Navigate to the SSJ directory and open the `rundirector` script in a text editor. Ensure the value of `set NUM_HOSTS` property equals the number of JVMs that will simultaneously connect to that controller. This will allow the director to know the total number of JVMs that will connect to the director, and allow the Director to wait until all connections are established before starting the benchmark.
 9. Please refer to Remote JVM Director
 10. Execute the `rundirector` script and wait until for a message stating that the director is ready for CCS connections. If not, something must not be synchronized, please refer back to check all steps were concluded.
 11. After that, change to the CCS directory. Execute `runCCS` to commence the benchmark.

Inventory of Operating System Services

After the `SPECpower_ssj2008` run is complete, the user is required to take an inventory of any services that were running during the benchmark run. After submitting results to SPEC, the user must retain this record for the duration of the review period as SPEC may request this information for results review purposes. Some suggestions for obtaining this information are:

- For Windows: `net start`
- For Red Hat Linux: `/sbin/runlevel; /sbin/chkconfig --list`

Results

When the benchmark iteration is finished, CCS calls the Reporter. The Reporter is another Java program whose function is to extract all of the benchmark data from the `.raw` file that CCS created. The Reporter formats this data into several human-readable HTML files. It will also use the power and performance data to create graphs for each JVM instance that was run, as well as for the overall `SPECpower_ssj2008` run in general. These graphs will be embedded into the HTML files. After the Reporter has finished creating the output files, all benchmark modules will terminate, and the user may then view the results of the benchmark. See section 6 for a detailed discussion of the `SPECpower_ssj2008` Results Reports.

The results in a multi-node configuration are stored differently than the configuration with one controller and only one SUT. Change to the results directory and locate the last run in the folder. This simply means this is the last run.

In this configuration, the run number and the machine name conveniently separates the output for each SUT, but the overall devices are precisely described in the original HTML output produced for the overall. So please make sure you update the configuration files on the SUT and the controller systems.

4 Operational Validity

In order to create compliant results, a run must pass several runtime validity checks. These checks are described in the *SPECpower_ssj2008 Run and Reporting Rules*.

5 Metric

5.1 SPECpower_ssj2008 Metric

The performance to power ratio of the target loads is a measurement of the number of ssj_ops (throughput) processed per watt of power consumed. The more ssj_ops the SUT can produce with one watt of power, the better the efficiency of the SUT.

The SPECpower_ssj2008 metric is calculated as the sum of all ssj_ops scores for all target loads, divided by the sum of all power consumption averages (in watts) for all target loads, including the active idle measurement interval.

$$\frac{\text{target load 1 ssj_ops} + \text{target load 2 ssj_ops} + \dots + \text{target load } n \text{ ssj_ops}}{\text{target load 1 watts} + \text{target load 2 watts} + \dots + \text{target load } n \text{ watts} + \text{active idle interval watts}}$$

The unit of the SPECpower_ssj2008 metric is “overall ssj_ops / watt”.

6 Results Reports

6.1 Overview

One of the major functions of CCS is to collect the benchmark data in real time, and output this data into reports that can be retrieved and viewed by the user. When a benchmark run is complete, there are a number of results files that CCS will have created for the user:

- a CSV file
- a RAW file
- several HTML files

Where these results files are stored can be specified by the user with the `ccs.results` property in the `ccs.props` file. (This property is not included in `ccs.props` by default, and must therefore be entered manually in the format `ccs.results=<desired results path>`.) By default, the results will be stored in `.../SPECpower_ssj2008/Results`. *The user is expected to ensure that the JVM instance for CCS will have the necessary credentials to write to this path.*

For each benchmark run, CCS will create a new directory under the Results directory into which it will store all of the results of the benchmark run. The name of the directory takes the form `ssj.<run number>`. For example, the first directory that CCS will create will be `ssj.0001`; the next directory (for the next run) will be `ssj.0002`, and so on.

This section will briefly explain the methodology employed by CCS for creating the results files. Details about each different type of results file will be discussed one at a time in the sections following this one.

When CCS is first initiated, one of the first things that it does is create two files: a CSV file and a RAW file. Throughout the entire benchmark iteration, CCS will collect data from the data sources to which it is connected, and store this data into these files in real time. The major difference between these two files is that the CSV file contains a sample-by-sample collection of measurement data, whereas the RAW file is used to store a more summarized version of the measurement samples, along with more details about each summary, and about the workload iteration.

Upon successful completion of the benchmark run, CCS will stop writing to the CSV and RAW files, and will then use the information contained in the RAW file to generate several human-readable HTML files. While all of the results files contain very valuable data about the benchmark run, many users will likely only be interested in the HTML files which are covered in detail in section 6.5.

6.2 The CSV file

The CSV (comma separated values) file is written in plain-text, Unicode format, and can be viewed with most popular text editors. The CSV file contains an extremely detailed collection of power consumption, temperature and performance data for the benchmark run at the highest possible level of resolution for each sample.

CSV File Naming Convention

The filename of the CSV results file takes the form of the directory name under which it resides, followed by `.ccs-log.csv`. For example, if the CSV file was created by CCS under the directory `.../SPECpower_ssj2008/Results/ssj.0008`, then the CSV filename would be `ssj.0008.ccs-log.csv`.

CSV File Format

The CSV file is divided into two main sections:

- a Test Bed Information Section
- a Measurement Section.

6.2.1.1 Test Bed Information Section

The Test Bed Information Section contains configuration information about CCS as well as detailed information about each data source that CCS is configured to connect to and collect data from.

```
#----- CCS properties -----
#Wed Nov 28 11:40:49 CST 2007
ptd.temp1.config.sensor.connectivity=USB
ptd.temp1.config.sensor.model=NS Thermal X
ptd.pwr1.config.ptd.os=Controller System
ccs.config.hw.cpu=Saturn ULP
... ..
... ..

#----- CCS System info -----
version=0.21
ccs.awt.toolkit=sun.awt.windows.WToolkit
ccs.file.encoding=Cp1252
ccs.file.encoding.pkg=sun.io
... ..
... ..

[wkld.MyWorkload] settings:
global.input.load_level.pre_measurement_seconds=30
```

```
global.input.log_level=INFO
global.os.name=Windows 2003
global.sun.io.unicode.encoding=UnicodeLittle
... ..
... ..

[ptd.MyPower] Model: Eureka Power Cyclone
[ptd.MyPower] Averaging interval(ms): 1000
[ptd.MyPower] . version=0.21-2f7044de

[ptd.MyTemp] Model: Acura NS Thermal X
[ptd.MyTemp] Averaging interval(ms): 5000
[ptd.MyTemp] . version=0.21-2f7044de

ptd.MyPower.ptd.version=0.21-2f7044de
ptd.MyPower.ptd.certified=True
ptd.MyTemp.ptd.version=0.21-2f7044de
ptd.MyTemp.ptd.certified=True
```

If CCS is configured to connect to more data sources, then information about those sources will be included here as well.

6.2.1.2 Measurement Section

After CCS records the test bed information into the CSV file, it then begins recording the measurement data that it retrieves from its data sources. The Measurement Section is arranged into a column-row format. How many columns actually get recorded to the CSV file will depend on the number of data sources that CCS has been configured to connect with. The column headings (for a minimal configuration) are:

- `CCS-Ser`: This column contains the serial number of the measurement sample, starting with 1, and on until the completion of the benchmark iteration.
- `Action`: This is the action requested by CCS from the workload and from the PTD instance(s). For each measurement sample, CCS will issue a "REQ_TX_DATA" request to its data sources; to which the data sources will respond with their most current measurement data. When the workload "state" changes (for example from target load 3 ramp-down to target load 4 ramp-up), CCS will post a "REQ_TX_SUMMARY" request to which the data sources will respond with a "summary" of their data for the workload interval. These actions determine the data that is returned and logged.
- `Run-ID`: This string allows use of a unique identifier for the benchmark iteration. It must be defined by the user with the `ccs.runId` property in the `ccs.props` file. It does not otherwise change.
- `wkld.MyWorkload.Trans`: The accumulated number of transactions that have been processed during the workload interval.
- `wkld.MyWorkload.BatchCount`: A Batch is a set of high level transactions that are executed during the workload iteration. It is the execution of these transactions that serve as the actual workload for the SUT.
- `pwr.MyPA.Watts`: The power reading, in watts, retrieved from Power & Temperature Daemon.
- `pwr.MyPA.Amp`: The current reading, in amperes, retrieved from Power & Temperature Daemon.
- `pwr.MyPA.Volt`: The voltage reading, in volts, retrieved from Power & Temperature Daemon.
- `pwr.MyPA.PF`: The power factor reading retrieved from Power & Temperature Daemon.

- `temp.MyTemp.Temperature`: The temperature reading as it is retrieved from Power & Temperature Daemon.
- `temp.MyTemp.Humidity`: The humidity reading as it is retrieved from Power & Temperature Daemon.

As more data sources are configured, more columns for those data sources will be included in the Measurement Section.

6.3 The RAW File

The RAW file is written in plain-text, Unicode format, and can be viewed with most popular text editors. In this regard, the RAW file only records the power, temperature and performance *summaries*. A summary is the average of all of the data samples that were captured during an interval. For example: if there were 240 samples of data captured during an interval, only the average of those 240 samples would be recorded into the RAW file. Also contained in the RAW file is a highly detailed collection of data about the workload iteration itself and about the actual hardware, software and benchmark configuration of the SUT(s).

The primary function of the RAW file is to serve as a source of data that can be parsed by the SPECpower_ssj2008 *reporter*. The reporter will be discussed in detail in section 6.4.

Raw File Format

The RAW file is ordered into 4 major sections.

- a test bed information section - is identical to the test bed information section of the CSV file
- a measurement section - the RAW file contains all of the same data types as those that are contained in the measurement section of the CSV file, along with many additional data types.
- a workload information section - contains very detailed statistics about each JVM instance (as well as the JVM Director) that was run within the workload.
- a global SUT information section - after the workload information section, the global SUT information section is recorded into the RAW file
- identification and contact information of the company that performed the benchmark run
- benchmark configuration information
- performance tuning information
- hardware information
- etc.
-

6.3.1.1 The Measurement Section

The data are organized in a line by line format, as opposed to a column-row format. Each line is of the format `field=value`. Every field is prefixed with the identifier `ssj2008.ccs.result.<serial number>`. The serial number is the consecutive number of the data sample for the RAW file, and is not to be mistaken for the serial numbers in the CSV file. Below is a (condensed) measurement sample from a RAW file:

```
ccs.result.28.CCS-ser=### [wkld.MyWorkload]
ccs.result.28.CCS_TIME=2007-06-19 10:19:32.484
ccs.result.28.ACTION=REQ_TX_SUMMARY
ccs.result.28.RUN_ID=MyRun
ccs.result.28.wkld.MyWorkload.CCS_RT=0
```

```
ccs.result.28.wkld.MyWorkload.WK_STATE=_04_lvl_sum_
ccs.result.28.wkld.MyWorkload.TRANS=-
ccs.result.28.wkld.MyWorkload.WHSE=-
ccs.result.28.wkld.MyWorkload.BATCH_RT=659179
ccs.result.28.wkld.MyWorkload.BATCH_COUNT=4107
ccs.result.28.wkld.MyWorkload.AVG_TXN=17113.033333333333
ccs.result.28.pwr.MyPA.CCS_RT=0
ccs.result.28.pwr.MyPA.WATT=310.059167
ccs.result.28.pwr.MyPA.AMP=2.663268
ccs.result.28.pwr.MyPA.VOLT=119.738708
ccs.result.28.pwr.MyPA.PF=0.979123
ccs.result.28.pwr.MyPA.NOTES="Watts,310.059167,223.700000,365.200000,
240,0,240,Volts,119.738708,119.170000,120.050000,240,0,240,Amps,2.663
268,1.934400,3.078000,240,0,240,PF,0.979123,0.948200,0.992000,240,0,2
40"
ccs.result.28.temp.MyTemp.TEMP=-1.0
ccs.result.28.temp.MyTemp.HUMID=-1.0
ccs.result.28.temp.MyTemp.NOTES="Temperature,-1.0,0,0,0,0,Humidity,-
1.0,0,0,0,0"
```

In this excerpt, the sample happens to be the 28th sample that is recorded into the RAW file.

6.3.1.2 The Global SUT Information Section

Each of these entries are prefixed with the identifier `ssj2008.wkld.ssj.global`. Below are some example entries:

```
ssj2008.wkld.ssj.global.user.country=US
ssj2008.wkld.ssj.global.user.dir=C:\SPECpower_ssj2008\ssj
ssj2008.wkld.ssj.global.user.home=C:\Documents and
Settings\Administrator
ssj2008.wkld.ssj.global.user.language=en
ssj2008.wkld.ssj.global.user.name=Administrator
etc...
```

6.4 The Reporter

Upon completion of the benchmark run, CCS will call the `SPECpower_ssj2008 reporter`, another java program whose function is to extract all of the benchmark data from the RAW file that CCS created. The reporter will format this data into a human-readable HTML file(s).

It will also graph the power and performance data in the RAW file that illustrates the throughput data for each JVM instance that was run. This graph illustrates the overall performance and power data (the final `SPECpower_ssj2008` score) for the entire benchmark implementation across all SUTs. The graphs will be embedded into the HTML file(s) by the reporter.

By default, the reporter will be run automatically by CCS as soon as the benchmark run is complete, creating the HTML file(s). However the reporter can be run manually, and doing so affords the opportunity for the user to take advantage of many command-line arguments and parameters that are supported by the reporter.

Running the Reporter Manually

Like the other modules of the `SPECpower_ssj2008` suite, the kit also comes with a run script that can be used to run the reporter. While the reporter can certainly be invoked by calling it manually with the java interpreter, the recommended (and easiest) way to

run the reporter is to use the run script included in the kit. The default name of the run script is `reporter.bat` for Windows, and `reporter.sh` for Unix.

The reporter run script is located in the `ssj` directory and is the only run script in the kit that requires command line arguments in order to function. The syntax for running the script is quite standard:

```
reporter.bat [options]
or
./reporter.sh [options]
```

as appropriate to the operating system. By initializing the reporter script it will dynamically display all the options afforded to the reporter. Additionally, the reporter can be used to compare the results of two benchmark runs. Below is an example command for comparing two RAW files:

```
./reporter.sh -e -r ../Results/ssj.0001/ssj.0001.raw -c
../Results/ssj.0002/ssj.0002.raw -o ../Results/compare.html
```

Results submitted to SPEC will appear on the SPEC website in the ASCII and HTML formats.

6.5 The HTML File(s)

Upon completion of the benchmark run, the reporter will generate an HTML results file with embedded graphs.

HTML File Naming Convention

There are 3 different types of HTML files that will be created by the reporter:

The filename of the overall HTML results file takes the form of the directory name under which it resides. For example, if the overall HTML file was created by CCS under the directory `.../SPECpower_ssj2008/Results/ssj.0008`, then the HTML filename would be `ssj.0008.html`.

Finally, the detailed HTML results files for each JVM instance takes the form of the directory under which it is stored, with `.details-<SUT_number>.<JVM_number>.html` appended to it. For example, `ssj.0008.details-0001.0003.html`.

Each type of HTML file is very easy to read, and therefore very self-explanatory. For this, only the overall HTML results file will be discussed in the section.

Overall HTML Results File Format

Unlike the other results files, the HTML files are designed to be highly “user-friendly” and very easy to read. Assuming the user has read and grasped most of the core concepts about the SPECpower_ssj2008 the user should have little to no difficulty reading and understanding the HTML results file, as it is extremely self-explanatory. For this reason, the HTML results file will only be discussed briefly.

The HTML file is made up seven major sections, with each section containing numerous subsets of data about the benchmark run.

The Header Section - It simply contains the most basic identification information about the benchmark the company running the benchmark, the platform that was tested, along with the SPECpower_ssj2008 score.

- Performance Results Summary - summarizes the power and performance data throughout the benchmark run, in addition displays performance to power ratios in easy-to-read table and graphical formats.
- System Under Test - configuration information about the SUT(s) including hardware, software, OS information, Java Runtime Environment and etc.
- Control and Collect System - information about the controller server's hardware and software, CCS, Java Runtime Environment information power analyzers and temperature sensors also.
- Tuning/Notes - discloses any performance tuning that was done within the test bed, along with any other miscellaneous notes that may be pertinent to the benchmark run.
- Electrical and Environmental Data - contains more details about the electrical data than what was disclosed in the Performance Results Summary section and environmental temperature data.
- JVM Instance Performance Section - contains the throughput performance data for each JVM instance that was run. The data is organized neatly into a table, and is also represented graphically within a benchmark iteration, the more entries there will be in this section

6.6 The SSJ Results Files

The user should be made aware that the SSJ module of the benchmark also generates its own results files.

The main purpose of the SSJ results files are for observing workload-specific information, as well as workload-specific debugging information. These results files can be found (by default) under the `.../SPECpower_ssj2008/ssj/results` directory of whichever system is hosting the JVM director.

The RAW Files

The RAW files for the SSJ module contain detailed information about the workload itself, and are used to generate the HTML results files for the SSJ module.

6.6.1.1 RAW File Naming Convention (Workload)

There are 3 types of RAW files that will be generated by the SSJ module:

The RAW file for the JVM Director takes the form `<directory_name>.director.raw`; where `directory_name` is the name of the directory in which the raw files are stored. For example: `ssj.0008.director.raw`.

The detailed RAW file for all JVM instances takes the form `<directory_name_.details.raw`. For example: `ssj.0008.details.raw`.

A detailed RAW file will also be created for each JVM instance. It takes the form `<directory_name>.<SUT_number>.<JVM_number>.raw`. For example: `ssj.0008.0001.0003.raw`.

The LOG File

The LOG file is a plain-text file and doesn't require much explanation at all. To put it briefly, it is a plain-text log to which the workload engine writes every time the workload performs any kind of action. Most users will probably only find it useful for occasional debugging purposes, as it contains time stamps for each time the workload performs a new action such as entering a new measurement interval. A log file is generated for each JVM instance.

6.6.1.2 LOG File Naming Convention

The LOG file takes the form `<directory_name>.<SUT_number>.<JVM_number>.log`. For example: `ssj.0008.0001.0003.log`.

The RESULTS File

There is one RESULTS file for each JVM instance as well as for the JVM Director. The RESULTS file is simply a recording of the STDOUT output of the JVM instance. In other words, all of the text that the workload engine displays to the console of the SUT(s) during the workload iteration is also copied into this file.

6.6.1.3 RESULTS File Naming Convention

The RESULTS file for the JVM Director takes the form `<directory_name>.director.results`. For example: `ssj.0008.director.results`.

The RESULTS file for each JVM instance takes the form `<directory_name>.<SUT_number>.<JVM_number>.results`. For example: `ssj.0008.0001.0003.results`.

The TXT File

The TXT file is simply an alternative format that is generated by the reporter, one for all JVM instances, and then one for each JVM instance. Each TXT file contains detailed information about the JVM instance.

6.6.1.4 TXT File Naming Convention

The TXT file for all JVM instances takes the form `<directory_name>.details.txt`. For example: `ssj.0008.details.txt`.

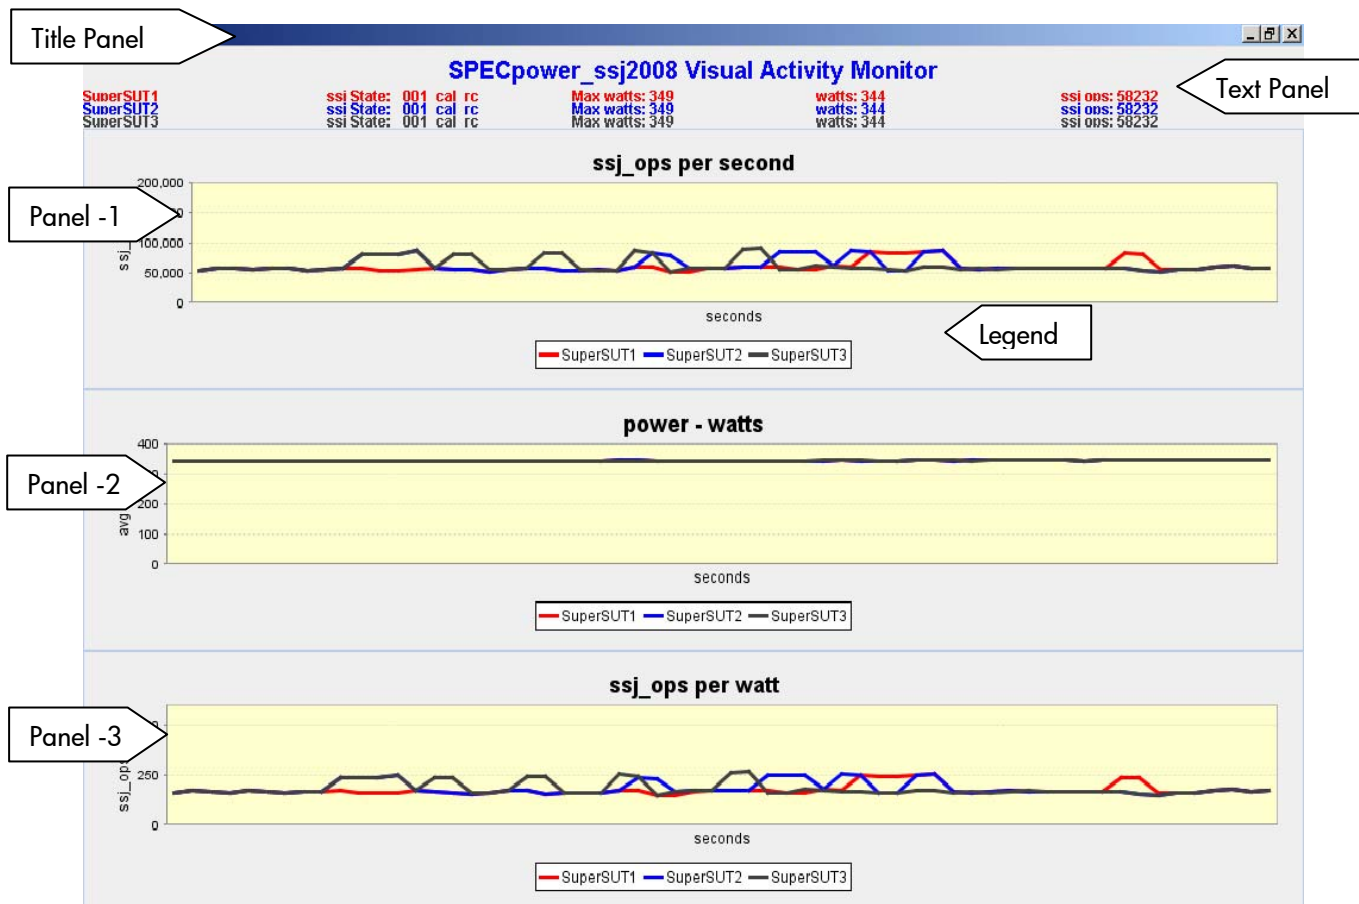
The TXT file for each JVM instance takes the form `<directory_name>.details-<SUT_number>.<JVM_number>.txt`. For example: `ssj.0008.0001.0003.txt`.

7 Visual Activity Monitor (VAM)

The Visual Activity Monitor (VAM) is an optional and flexible Java-based software tool designed to render a real-time graphical representation based on data collected during the SPECpower_ssj2008 benchmark execution. It is capable of retrieving performance data from up to three (3) SUTs, as it is captured by the CCS. The VAM can run in virtually any operating environment using configuration and communication methods similar to those employed by other SPECpower benchmark software components. The VAM generates a compelling visual representation of any performance impact produced by the introduction of new server technologies. Communication occurs between predefined

TCP/IP ports associated with one or more instance of CCS. The VAM display (Figure 7.2) is organized into 'panels,' each corresponding to a set of entries located within the control file 'vam.props.' Panels can be omitted by removing or simply commenting the associated entries within the 'vam.props' file.

Figure 7.2 – Sample VAM Display Window with Legend



Installing VAM

No additional installation is required. VAM is distributed and installed along with all the other components of SPECpower_ssj2008 1.10. The VAM code, batch file (runvam.bat) and property file (vam.props) are located in the VAM subdirectory of SPECpower_ssj2008. Ensure that the runvam.bat is updated with the path to Java in order to work correctly.

```
set JAVA=c:\yourjava\bin\java
set SSJ_HOME="..\ssj"
:: the line below establishes the path to the required libraries
set CP%SSJ_HOME%\lib\jfreechart-1.0.6.jar;%SSJ_HOME%\lib\jcommon-1.0.10.jar
%JAVA% -classpath %CP% vam.VAMApp vam_x.0.0.props
```

Verify or Establish Required Communication

Verify that the benchmark is behaving as expected. Check for network accessibility between servers. If necessary, use ping to verify connectivity. All servers must exist on a common private subnet. In addition, all TCP port assignments associated with the requisite software and hardware components should be properly coordinated before proceeding. If problems are encountered, review all preceding sections that may apply.

VAM Setup

Preparing VAM for use requires changes to the properties file 'vam.props' to include the desired panels. A unique user-defined name has to be assigned to each SUT. In addition, each CCS instance (i.e. - ccs.ptd.pwr1) that is expected to provide data to VAM, must include entries in the 'ccs.props' file to "turn on" the TCP port through which the corresponding CCS instance will communicate with VAM. The appropriate declaration (i.e. - ccs.vam =vam1) within the 'vam.props' file needs to be included.

When benchmarking multiple sources, each JVM instance is controlled and coordinated by the ssj_2008 Director. The Director consolidates the performance data from multiple JVM instances and passes an aggregate ssj_ops per second to the CCS which is also gathering data from multiple power analyzers. If configured for multiple sources, all data provided by CCS and presented by the VAM display window represent the aggregate of each value. Set the 'syncstart' property to "true" in the vam.props and ccs.props files in order to synchronize the data flow between VAM, CSS and SUTs.

The following example represents the maximum number of data sources and describes changes required to introduce VAM to a functioning benchmark

vam.props

```
vam.datasources = sut1, sut2, sut3 (assigns user-defined device names)

sut1.name = SuperSUT1           (name reflected on chart)
sut1.color = red                (representative color for 'sut1')
sut1.port = 8908                (uniquely assigned TCP port)

sut2.name = SuperSUT2
sut2.color = blue
sut2.port = 8909

sut3.name = SuperSUT3
sut3.color = darkgray
sut3.port = 8910

vam.syncStart = true
```

css.props

```
css.vam = vam1, vam2, vam3      (declares data sources 'vam1', 'vam2', 'vam3')

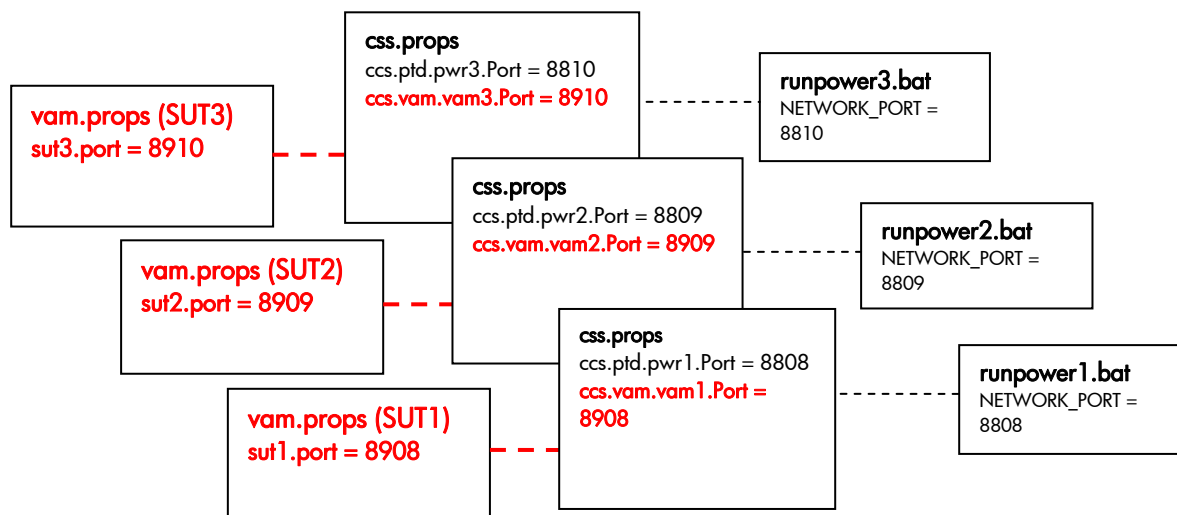
css.vam.vam1.type = VAM         (identifies data source as VAM)
css.vam.vam1.IP = localhost     (location of VAM instance)
css.vam.vam1.Port = 8908       (unique TCP/IP Port for VAM instance)

css.vam.vam2.type = VAM
css.vam.vam2.IP = localhost
css.vam.vam2.Port = 8909

css.vam.vam3.type = VAM
css.vam.vam3.IP = localhost
css.vam.vam3.Port = 8910

css.vam.syncstart = true       (synchronizes VAM, CCS and SSJ instances)
```

Figure 7.3 – Describes TCP port coordination with multiple SUTs



VAM Presentation Properties

An experienced user may choose to tailor the display with a particular audience in mind. Therefore many properties are available within the 'vam.props' file that can be used to customize the appearance of the VAM display window.

Global Settings

The VAM display window refresh rate is expressed in seconds.

```
vam.refreshRate = 1
```

Adjust the number of data points represented on the X axis in coordination with the smoothing feature for optimum data presentation. Default smoothCount value is 1.

```
vam.chartDataPoints = 60
```

```
vam.smoothCount = 2
```

Screen Size and Location

VAM window starting position and size

X = horizontal window start position, e.g. 100 pixels from left side of display

Y = vertical window start position, e.g. 120 pixels from top of display

Width and Height represent frame dimensions

```
vam.frameTopX = 100
```

```
vam.frameTopY = 120
```

```
vam.frameWidth = 800
```

```
vam.frameHeight = 900
```

Title Panel and Text

This property set defines the appearance of the Title Panel located across the top of the VAM display window.

`title.panel.show` configures VAM to display the title panel.

```
vam.titlePanel.show = yes
```

`titlePanel.title` is used to title the VAM display window.

```
vam.titlePanel.title = SPECpower Visual Activity Monitor
```

`titlePanel.fontcolor` is used set the Title's text color.

```
vam.titlePanel.fontColor = blue
titlePanel.fontcolor is used set the Title's text font size.
vam.titlePanel.fontSize = 20
"size.xMin" determines fixed title box width.
vam.titlePanel.size.xMin = 100
"size.yMin" determines fixed title box height.
vam.titlePanel.size.yMin = 30
```

Chart Panel Example

Setup for the watts chart

vam.wattPanel.show configures VAM to display the Power-Watts panel.

```
vam.wattPanel.show = yes
```

vam.wattPanel.title provides the panel title.

```
vam.wattPanel.title = power - watts
```

vam.wattPanel.x-axis labels the x-axis

```
vam.wattPanel.x-axis.title = seconds
```

vam.wattPanel.y-axis labels the y-axis

```
vam.wattPanel.y-axis.title = avg watts
```

vam.wattPanel.legend displays the legend for each panel of set to "yes."

```
vam.wattPanel.legend = no
```

y range will default to auto when value is blank

```
vam.wattPanel.y-Value.min = 0
```

```
vam.wattPanel.y-Value.max = 50
```

Adjusting Panel Size

Default size is 100, 50 if values are blank

```
vam.wattPanel.size.xMin = 100
```

```
vam.wattPanel.size.yMin = 100
```

Freeze Button Example

The "freeze button" will stop the data display, essentially freezing the lines on the charts.

This capability is useful for demos to explain some specific behavior. Toggle between "Freeze/Resume" to utilize this feature. Data points produced during the freeze period cannot be displayed.

Display the "Freeze" button by setting buttonPanel.show to "yes."

```
vam.buttonPanel.show = yes
```

```
vam.buttonPanel.size.xMin = 100
```

```
vam.buttonPanel.size.yMin = 30
```

Colors

The list below shows the only valid entries where 'color' can be changed.

Valid Colors List:

Black, Blue, Cyan, Darkgray, Gray, Green, Lightgray, Magenta, Orange, Pink, Red, White, Yellow

Starting VAM

Proceed as if starting SPECpower_ssj2008 in accordance with preceding instruction. After the director has acknowledged each of the workloads (SUTs) and PTD devices execute the CSS (runccs.bat) script and start VAM via the runvam.bat script. VAM will begin presenting the incoming data immediately. If the 'syncstart' property is "true," VAM and CSS will each wait for the other before data is presented to the VAM display window.

Caveats

With release 1.0, VAM cannot be stopped and then re-started with the same instantiation of CCS, VAM will not “re-connect”.

A given instance of CCS can only provide data to one instance of VAM representing a maximum of three (3) SUTs.

Figure 7.6 – Shows benchmark with VAM included

```

Time,02-10-2009 12:49:20.437,Watts,364.000000,Volts,
ark,_001_cal_ru_
Response to client: Watts,364.000000,Volts,119.110000

Time,02-10-2009 12:49:21.453,Watts,367.280000,Volts,
ark,_001_cal_ru_
Response to client: Watts,367.280000,Volts,119.170000

C:\SPECpower_ssj2008\vam>java -classpath vam.jar;..\ssj\lib\jfreechart
ib\jcommon-1.0.13.jar vam.VAMApp
*****
vam.jar version 1.0.0   December 17, 2008
Licensed Materials - Property of SPEC
Copyright 2006-2008 Standard Performance Evaluation Corporation (SPEC)
All Rights Reserved
*****
SuperSUT1 is waiting for connection
SuperSUT2 is waiting for connection
SuperSUT2 is connected
SuperSUT1 is connected

Time,02-10-2009 12:49:20.515,Watts,347.750000,Volts,
ark,_001_cal_ru_
Response to client: Watts,347.750000,Volts,118.970000

Time,02-10-2009 12:49:21.500,Watts,349.610000,Volts,
ark,_001_cal_ru_
Response to client: Watts,349.610000,Volts,119.020000

Time,02-10-2009 12:49:18.187,Temp,21.222000,Humidity,5
Response to client: Temperature,21.222000,Humidity,5

Time,02-10-2009 12:49:19.187,Temp,21.222000,Humidity,5
Response to client: Temperature,21.222000,Humidity,5

Time,02-10-2009 12:49:20.187,Temp,21.222000,Humidity,5
Response to client: Temperature,21.222000,Humidity,5

Time,02-10-2009 12:49:21.187,Temp,21.222000,Humidity,5
Response to client: Temperature,21.222000,Humidity,5

vam.vam2: Received(12:48:36): get
vam.vam1: Received(12:48:36): get
vam.vam2: Received(12:48:37): (null)
vam.vam1: Received(12:48:37): (null)
vam.vam1: Received(12:48:37): _ready_
vam.vam2: Received(12:48:37): _ready_
====> Ready for testing <====
start: wkld.ssj_dir
start: ptd.pwr1
start: ptd.pwr2
start: ptd.temp1
====> Start testing <====

Uncertainty of clock offsets is ~ 2ms
Waiting for jvm "DST0-SUT2.001" data collector on port 1500
Accepted data collector client for "DST0-SUT2.001" at /10.1.1.103:1040
Waiting for jvm "cp-ctrl1.001" data collector on port 1500
Accepted data collector client for "cp-ctrl1.001" at /10.1.1.101:1043
Director ready for CCS connection
jar validity is true
Directing 2 jvm instances on 2 hosts.

Preparing to synchronize for 3 calibration measurements...
interval 1 : Reading READY state from the jvm instances
interval 1 : Received READY from all, now sending START
interval 1 : Calibration 1 run in progress...

```

Quickly Customizing Display

To quickly customize the display, VAM can run without any instances of CCS.

Essentially, you will iterate through the following sequence of actions until you are satisfied.

1. open the VAM properties file for editing
2. make the necessary changes and “save the file”
3. run VAM, see changes required.
4. exit VAM
5. edit the already open VAM properties
6. Repeat (go to step 2)

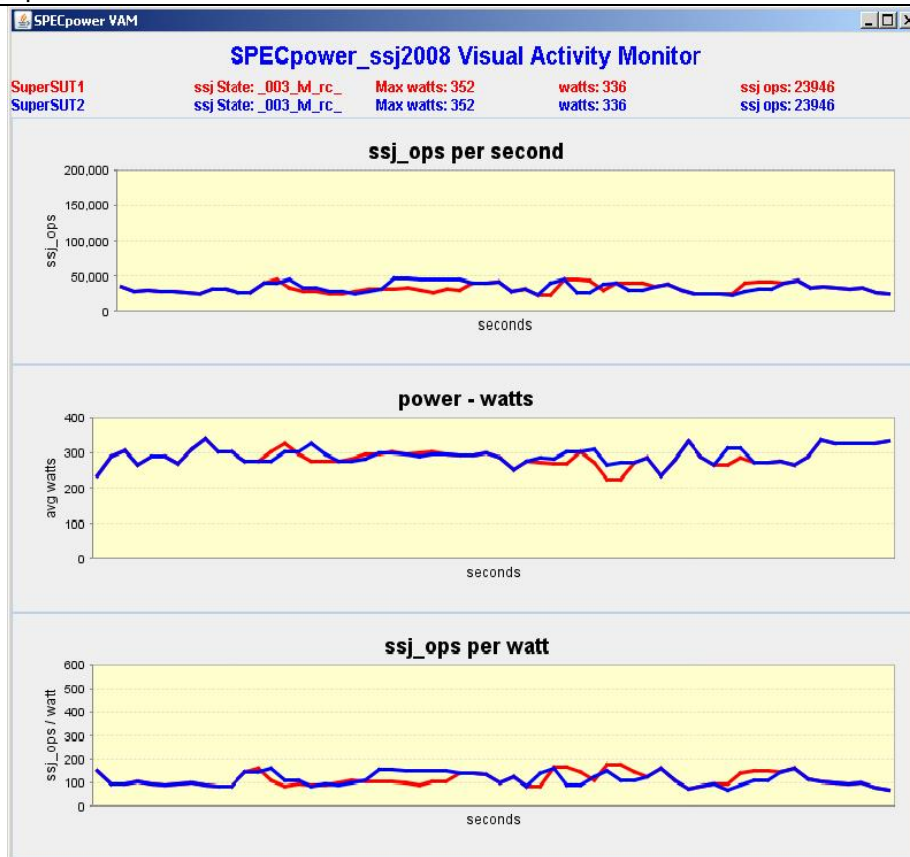


Figure 7.5 – Customized Display

8 Performance Tuning

All of the techniques and concepts that are involved in this area reach indefinitely beyond the scope of this document. Therefore, only a few elementary concepts that are related to the SPECpower_ssj2008 benchmark will be briefly discussed in this section. This document is not to be considered, and does not claim to be, a complete reference for SPECpower_ssj2008 performance tuning.

Please note that these techniques may or may not necessarily be beneficial for system performance in a production environment.

- **JVM Software:** This is a critical component affecting the performance of the system for the workload. Use JVM profiling (java -prof or java -Xprof) to identify the most heavily called API routines. Other tools may be available on your system.
- **CLASSPATH:** Having extra items in CLASSPATH can degrade performance on some JVMs.
- **Memory:** The SPECpower_ssj2008 workload uses at least a heap size of 256 and is known to benefit from additional memory. Heap size is set via the -ms/-mx (or -Xms/-Xmx) command line arguments for the java executable in many JVMs.
- **Threads:** The number of threads is approximately equal to the number of warehouses. For very large numbers of warehouses, you may need to increase operating system limits on threads.
- **JVM Locking:** At a larger number of warehouses (and hence threads), the synchronization of application methods, or the JVM's internal use of locking, may prevent the JVM from using all available CPU cycles. If there is a bottleneck

without consuming 100% of the CPU, lock profiling with JVM or operating system tools may be helpful.

- **Network:** Although SPECpower_ssj2008 uses very little network I/O, it is a good idea to make sure that the network fabric being used for the test bed has a very high availability for use by SPECpower_ssj2008. The best practice is to setup the test bed on its own isolated broadcast domain. This will minimize the possibility for network anomalies that may interfere with SPECpower_ssj2008's network-specific components.
- **Disk I/O:** SPECpower_ssj2008 does not write to disk as part of the measured workload. Classes are read from disk, of course, but that should be a negligible part of the workload. If the disk I/O is found to be higher than it should be, then there may be another process accessing the disk during the benchmark run. While this may or may not be affecting performance, the excessive disk I/O may be contributing to higher power consumption, which will ultimately lower the SUT's final score.

Unlike most benchmark implementations, in SPECpower_ssj2008, tuning for performance is not always the most important consideration for the user. The user should also be mindful of the SUT's power consumption characteristics. Remember, a higher power consumption will lower the SPECpower_ssj2008 score.

As with most aspects of the benchmark implementation, there are restrictions concerning what kind and how much performance and power-consumption tuning is allowed for publishable results. Please see the *SPECpower_ssj2008 Run and Reporting Rules* for details.

9 Submitting Results

Upon completion of a compliant run, the results may be submitted to SPEC. Please see the *SPECpower_ssj2008 Run and Reporting Rules* for details.