

SZSN Expansion To Increase Production By 60 Million KG

Shandong Zhouyuan Seed and Nursery Co., Ltd (SZSN)

\$0.28 UP 16.6%

SZSN's 45,000 square foot expansion and new advanced production line will increase annual production by 60 Million KG's of seed.

This one is on the move again! Watch for news and get on SZSN first thing Monday!

Other java methods may be dynamically compiled using HotSpot, and are referred to as compiled methods.

The time has arrived for today's developers!

The second reason why profile feedback is not more widely used is that there is a concern that a performance gain for one workload may be at the expense of the performance of another workload.

sh script in the WebLogic installation, and use it as a template for collectlaunch.

The simplest case of deadlock occurs when two threads hold locks for different r

resources and each thread tries to acquire the lock for the resource held by the other thread.

Using other specifiers it is possible to record other data, such as the effective address of memory operations.

Some of the JVM frames represent transition code between interpreted Java, compiled Java, and native code.

When looking at the Timeline in Analyzer, you may want to color methods from the WebLogic infrastructure all one color, and use other coloring to look at specific sections of your code.

In this article, a simple and popular parallel partitioning program example is used to illustrate the above various race condition issues.

Using other specifiers it is possible to record other data, such as the effective address of memory operations.

In addition, other code is generated in the data space to execute the transitions between interpreted and compiled code.

Java programs may have explicit synchronization, usually performed by calling the monitor-enter routine.

But data race situations are generally very difficult to find.

Therefore it does not make sense to try and compare the absolute count numbers.

Running the collect command with no arguments will run a script that will verify that all the appropriate patches are installed.

The stop command will terminate the profiling run.

If they fail to do so, a data race could result, causing the program to produce different results when rerun with the same input.

java file from which it was compiled, with metrics on each source line.

Source code annotations are assertions and NOTES that you place in your source code to pass information to LockLint.

Threads must acquire and release locks associated with the data they share.

sh script or the startManagedWebLogic.

This routine takes a list of the instruction types to trace.

The next step is to select which instructions are to be traced, and what is to be recorded about those instructions.

VOLATILE imposes limits on optimization.

The objective for the analyzer being written here is to print out the traced instructions.

For more information, see The Java Technology Brand.

The machine representation looks the same as that for any other application in a traditional language.

Among the operations that can be specified is the dynamic loading of a shared object, and calls into various functions or methods contained in that object.

The overhead of these routines, and the cost of writing the experiments to disk will dilate the runtime of the Java program.

c to write a rule to create a .

Unknown side effects can lead to data races or deadlocks.

If you want to examine it on a different machine, you should set -A copy to make the experiment more portable.

Where Do I Go To Get Help?

The default uses a period to separate the whole number and decimal parts of a floating-point number formatted with D, E, EN, ES, F, and G editing.

This is particularly true for vars and funcs, which put out a single line of information for each variable or function.

The next step is to select which instructions are to be traced, and what is to be recorded about those instructions.

so and pass it a Java class.

Most processes share several resources.

The default experiment is a clock-profiling experiment.

What is a Sun Studio Express Build?

It specifies that the value of an entity might be changed or referenced by means other than Fortran.

You can also use the -d argument to specify a directory into which the experiment

ts should be stored.

In the partitioning example shown below, the main program processes the input source data into an object array and then creates N working threads to sort and collect the objects in parallel.

Most software developers are trained to write their code as a single flow of sequential operations.

The methods presented here are necessary conditions for the training workload to be representative of the reference workload.

To present the data visually, the basic blocks are ordered according to their counts, basic blocks with higher counts are plotted further from the origin.

If there are multiple training workloads, the aggregate data for all of them should be gathered.

A good example of what can be achieved is to run the BIT tool on the test program shown below.

If different operations use a common set of resources, but the order in which they acquire the locks is inconsistent, there is a potential for deadlock.

To profile a server, you must ensure that the JVM command launching the server is prepended with a collect command, with appropriate arguments, to invoke the Sun Studio Collector.

To profile a server, you must ensure that the JVM command launching the server is prepended with a collect command, with appropriate arguments, to invoke the Sun Studio Collector.

This information can point to a potential data race.

The next step is to select which instructions are to be traced, and what is to be recorded about those instructions.

Such use does not interfere with SIGPROF usage for clock-profiling.

The overhead of these routines, and the cost of writing the experiments to disk will dilate the runtime of the Java program.

Operations within the application may require access to more than one of those resources.

This new OpenSolaris-based distribution provides an integrated environment for the development of Solaris, Java, and AMP-based applications.

One of these is the actual interpreter code used to process the applications Java byte-code methods.

Rather memory is managed by a so-called garbage-collector.

Source code annotations are assertions and NOTES that you place in your source code to pass information to LockLint.

Submit Your Tips and Get Cool Gifts Send in your articles and tech tips on Solaris software development.

Methodology The methodology proposed here is to gather both basic block count and branch probability data for the training and reference workload.

Unknown side effects can lead to data races or deadlocks.

Similarly assume that R_i is the number of times that basic block i is executed during the run with the reference workload.

Coverage can be estimated from basic code block counts.

The methods presented here are necessary conditions for the training workload to be representative of the reference workload.

In the partitioning example shown below, the main program processes the input source data into an object array and then creates N working threads to sort and collect the objects in parallel.

The information that is most useful for determining whether the training workload is appropriate are the basic block counts, and the branch taken probabilities.

Excelsior, LLC Fujitsu Gemstone Systems, Inc.

Additionally, for parallel programs written in Sun Fortran, it can be used to specify that the value of an entity might be changed by another thread than the current thread.

The Analyzer has a radio button in the Data Presentation Dialog for turning view mode to user, expert, or machine.

Server Profiles This section describes tips for examining the server profiles.

You may also use hardware counter profiling or synchronization tracing. You may make assertions about which variables are supposed to be protected by a lock and which locks are supposed to be held whenever a function is called. Additionally, for parallel programs written in Sun Fortran, it can be used to specify that the value of an entity might be changed by another thread than the current thread.

The following is a step by step walkthrough of the process of gathering the data for either the training or reference workload.

This new OpenSolaris-based distribution provides an integrated environment for the development of Solaris, Java, and AMP-based applications.

If you want to examine it on a different machine, you should set -A copy to make the experiment more portable.

The information that is most useful for determining whether the training workload is appropriate are the basic block counts, and the branch taken probabilities.

Core Dump Management on the Solaris OS Find out about features and analysis tools in the Solaris OS that can be used to manage core dumps.

Descendant Process Controls Many WebLogic Servers do not spawn additional processes, so -F on is not needed.

Some of the JVM frames represent transition code between interpreted Java, compiled Java, and native code.

Whilst coverage is a minimum criteria for acceptance of a training workload, having good coverage does not necessarily mean that the training workload is ideal.

Using kstat From Within a Program in the Solaris OS Find out why you should use kstat and libkstat to extract metrics from the Solaris OS.

h header file which can hold information about the values held in the source and destination registers.

The following tables list the routines of the Solaris OS and POSIX libthread APIs recognized by LockLint.

ll files, first use the rule for creating a .

But data race situations are generally very difficult to find.

For clock- and HWC-profiling, the data collection process makes various calls in to the JVM, and handles profiling events in signal handlers.

With 'COMPATIBLE', the value resulting from data conversion is the one closer to the two nearest representations, or the value away from zero if the value is halfway between them.

The following example demonstrates the use of the VALUE with a C program calling a Fortran subprogram with a literal constant as an argument.

This information enables the compiler to make better optimisation decisions.

Join us at a city near you.

Some data race conditions are not harmful and can be permitted in parallel software for performance reasons.

ll files, first use the rule for creating a .

Product and Technology Releases Solaris Express, Developer Edition Check out this free integrated Solaris development environment.

It is actually necessary to have a bit more complexity than this to handle the forking of the traced application.

The methods presented here are necessary conditions for the training workload to be representative of the reference workload.

However, there is not always a one-to-one correspondence between the two.

Protecting CIFS Clients From Viruses Using Sun StorageTek NAS OS Need to protect your Common Internet File System clients?

The next step is to select which instructions are to be traced, and what is to be recorded about those instructions.

Using kstat From Within a Program in the Solaris OS Find out why you should use kstat and libkstat to extract metrics from the Solaris OS.

Both callstacks are recorded during profiling, and are reconciled during analysis.

The information that is most useful for determining whether the training workload

d is appropriate are the basic block counts, and the branch taken probabilities.

Typically doing two passes through the compiler will, unavoidably, take about twice as long as doing a single pass.

But data race situations are generally very difficult to find.

Methodology The methodology proposed here is to gather both basic block count and branch probability data for the training and reference workload.

The most often encountered race conditions are data race conditions.

It will not have any frames corresponding to Java housekeeping code, which does not have a Java representation.

During execution of a Java application, most Java methods are interpreted by the JVM; these methods are referred to below as interpreted methods.

Threads must acquire and release locks associated with the data they share.

PROTECTED imposes limitations on the usage of module entities.

There are three modes of looking at the data from Java programs: the user representation, the expert representation, and the machine representation.

The -xbinopt flag tells the compiler to produce an annotated application, suitable for further analysis at a later time.

Methodology The methodology proposed here is to gather both basic block count and branch probability data for the training and reference workload.

A good example of what can be achieved is to run the BIT tool on the test program shown below.

Rather memory is managed by a so-called garbage-collector.

This means you and your customers can be assured that these applications are portable across the products available today that are compatible with Java SE technology.

Caller-callee relationships will show all overhead frames, and all frames representing the transitions between interpreted, compiled, and native methods.

Then, set properties for the Node Manager to tell it to use the new script, and restart the Node Manager.

The Timeline in the Java representation shows all Java threads.

The Timeline in the machine representation will show bars for all threads, LWPs, or CPUs, and the callstack in each will be the machine-representation callstack.

Since WebLogic servers are Java-based, the -j on option to collect is always needed.

For example, which routines are appropriate to inline, or which branches are the frequently taken path.

LockLint subcommands help you analyze your code and discover which variables are not consistently protected by locks.

A data race condition is caused by unordered concurrent accesses of the same memory location from multiple threads.

The methods presented here are necessary conditions for the training workload to be representative of the reference workload.

There are some situations where race conditions in a parallel program for performance reason.

The most often encountered race conditions are data race conditions.

You might find that your problem has already been solved.

In Part I, he discusses the principles and illustrates them with an example that works only with password-based authentication.

On Solaris, the Node Manager, by default, uses native libraries to do the launch, but it can be told to use a startup script instead.

Users should allocate a character buffer large enough to hold the longest message.

Java profiling is not yet supported for this usage.

For more information, see The Java Technology Brand.

The second reason why profile feedback is not more widely used is that there is a concern that a performance gain for one workload may be at the expense of the performance of another workload.

o file is produced when you compile with the -Zll flag.

The API was primarily designed to support statistical sampling of callstacks, but it also contains extensions to support tracing.

Race conditions are hard to avoid and also hard to find with the conventional debugging methods and tools.

This paper presents two ways of viewing the correspondence between the behaviour of the training and reference workloads.

This can assist in judging the appropriateness of the chosen granularity, that is, which variables are protected by which locks.

Less frequent but harder to find are general race conditions.

Java profiling is not yet supported for this usage.

Solaris Express, Developer Edition Check out this free integrated Solaris development environment.

Furthermore a race condition might be the symptom of a deeper design problem in the program code.

This can assist in judging the appropriateness of the chosen granularity, that is, which variables are protected by which locks.

Source code annotations are assertions and NOTES that you place in your source code to pass information to LockLint.

Therefore it does not make sense to try and compare the absolute count numbers.

The current value of the argument is passed to the subprogram and cannot be changed in the subprogram.

The function list in the Java representation shows metrics against the Java methods, and any native methods called.

Protecting CIFS Clients From Viruses Using Sun StorageTek NAS OS Need to protect your Common Internet File System clients?

Only Java SE licensees can claim compatibility with Java SE specifications and can ship Java SE-branded products.

Rather memory is managed by a so-called garbage-collector.

Core Dump Management on the Solaris OS Find out about features and analysis tools in the Solaris OS that can be used to manage core dumps.

The following command will run an instrumented application and output basic block counts and branch probabilities.

You can make a number of useful queries about the functions, variables, function pointers, and locks in your program.

cshrc source the file in your current working directory, since others may want to run LockLint on those same files, and they may not use the same shell you do.

Most processes share several resources.

The main thread is used for Java initialization, and then becomes the user's java-program main thread; additional user threads are created as needed.

To present the data visually, the basic blocks are ordered according to their counts, basic blocks with higher counts are plotted further from the origin.

More details on binary optimisation with -binopt can be found in a recent article.

An example of such a complex script appears below.

Java always uses bound threads, so that there is a one-to-one correspondence between threads and LWPs.

It is actually necessary to have a bit more complexity than this to handle the forking of the traced application.

Makefile Rules To modify your makefile to produce .

so and pass it a Java class.

c The man page for BIT describes the options that are available for the tool.

Variables that were protected by a given lock.

If different operations use a common set of resources, but the order in which they acquire the locks is inconsistent, there is a potential for deadlock.

The caller-callee panel shows the calling relationships in the Java representation.

Therefore it does not make sense to try and compare the absolute count numbers.

Both machine- and Java callstacks are collected for these events, but no synchronization tracing data is collected for internal locks used within the JVM.

A workaround to support profiling for such an application, is to not set -j on f

or the initial collect invocation, and to ensure that the -Xruncollector option is passed to the invoked JVM.

During execution of a Java application, most Java methods are interpreted by the JVM; these methods are referred to below as interpreted methods.

The experiment begins when the server is launched, and terminates when the server exits.

In the current release, clock- and hw-counter-profiling, as well as synchronization tracing are supported.

You would send the same signal to pause data collection either shortly before or just after the benchmark load is terminated.

In the current release, clock- and hw-counter-profiling, as well as synchronization tracing are supported.

Both callstacks are recorded during profiling, and are reconciled during analysis.

The current value of the argument is passed to the subprogram and cannot be changed in the subprogram.

Alternatively, you can use LockLint subcommands to load the relevant .

Source for a Java method corresponds to the source code in the .

Dynamically compiled methods are loaded into the data space of the application, and may be unloaded later.

sh; the two properties to be set are: Property Value StartTemplate collectlaunch.

The most challenging issue for developers new to the parallel software paradigm is a fundamental problem in parallel programming: syndromes known as race conditions.

Typically doing two passes through the compiler will, unavoidably, take about twice as long as doing a single pass.

Rather memory is managed by a so-called garbage-collector.

Threads must acquire and release locks associated with the data they share.

In this case, we choose to name the script collectlaunch.

In addition, other code is generated in the data space to execute the transitions between interpreted and compiled code.

The next step is to select which instructions are to be traced, and what is to be recorded about those instructions.

Therefore it does not make sense to try and compare the absolute count numbers.

In the machine representation, multiple HotSpot compilations of a given method will be shown as completely independent functions, although the functions will all have the same name.

Coverage can be estimated from basic code block counts.

Core Dump Management on the Solaris OS Find out about features and analysis tools in the Solaris OS that can be used to manage core dumps.

Then, set properties for the Node Manager to tell it to use the new script, and restart the Node Manager.

The information that is most useful for determining whether the training workload is appropriate are the basic block counts, and the branch taken probabilities.

However, there is not always a one-to-one correspondence between the two.

The callstack for a Java thread may have a mixture of Java PCs and machine PCs.

You may want to use the date command to generate a string representing the current date and time and embed that string in the experiment name.

VOLATILE imposes limits on optimization.

Most programmers report that they find source code annotations preferable to command-line subcommands.

Be aware that using any of the features may make your program source code difficult to port to other platforms and compilers that do not yet support these features.

Accesses to variables that are not consistently protected by at least one lock, and accesses that violate assertions about which locks protect them.

Source from compiled methods will be shown against the Java source; the data will represent the specific instance of the compiled-method selected.

There are three modes of looking at the data from Java programs: the user representation, the expert representation, and the machine representation.

In the current release, clock- and hw-counter-profiling, as well as synchronization tracing are supported.

Source for a Java method corresponds to the source code in the .

BIT can gather data on the number of times basic blocks are executed, or the probability that a given branch is taken.

Join us at a city near you.

Initiating Data Collection This section describes how you can modify the scripts used to launch WebLogic servers to enable data collection.

Collecting the Data Profiling a Java application is done by collecting data on the JVM as it runs the Java application.

A Java program is inherently multithreaded, and has one JVM thread for each thread in the user's Java program.

Node-Manager Launched Servers The BEA Node Manager is a Java process that controls servers, based on requests received from the Administration Console.