





Trace Genera

Overview

Conceptually, *J is a pipeline consisting of three major components. The trace generator produces and serializes a stream of events from a program execution in the form of a trace file. The trace is then passed on to the trace analyzer, which is primarily used to compute dynamic metrics. Finally, the metric data is emitted in the form of an XML file, which is then ready to be processed and incorporated into a database constituting the core of a dynamic website, or viewed using other means.

*J is distributed at: www.sable.mcgill.ca/~bdufou1/starj/

Traces in *J are formed of records of runtime events, and use a simple, generic and task of collecting data from multiple sources. This design also makes it easy to add by reprocessing old data, thus avoiding costly trace generation phases.

Currently, our trace generator uses the buit-in Java Virtual Machine Profiler Interf from an executing Java Virtual Machine (JVM). Command-line options can mod specification file can be used to select which events and which of their fields have requiring any modification to the agent. This agent runs on any platform impleme

Unfortunately, while the JVMPI is reasonably ubiquitous it has several drawbacks JVMPI's callback mechanism is inherently slow for frequent events, which results time even for simple benchmark programs. The JVMPI implementation also does events may be skipped during VM startup and have to be explicitly requested in o requires non-trivial internal state in order to track and handle missing events. A fu it provides. Not only are the event types fixed (limiting dynamic metric possibilitie the JVMPI reports instruction executions using code offsets, and so locating the ac requires classfile parsing.

Complete traces of even small programs can be very large. *J thus supports splitting *J is also designed to make use of a pipe or socket between the agent and the anal files on disk. While *J can output GZipped trace files, more appropriate compress generators are part of active future work.

*J: A Tool for Dynamic Analysis of Java Programs

Requests Instruction Predictor ead ID Set nod ID Set na ID Set	<section-header><section-header><list-item><list-item><list-item><list-item><list-item><list-item></list-item></list-item></list-item></list-item></list-item></list-item></section-header></section-header>	Trace Instruction Predictor
ator at extensible trace format. This simplifies the dd new metrics because they can be calculated rface (JVMPI) to dynamically receive events lify the agent's behaviour, and in particular, a e to be recorded in the trace file, without enting the JVMPI specification. The most obvious one is the fact that the ts in a very significant increase in execution s not guarantee complete trace data: some order to obtain a precise execution trace. This further limitation of the JVMPI is in the data ies), but even the event data can be insufficient: actual opcode of the executed instruction ing trace files into more manageable sizes. lyzer, eliminating the need to store trace sion techniques and alternative trace	<text><text></text></text>	The analyzer in *J processes the s modular metric computations and to the set of analyses. Analysis operations in *J are orga whereas Operations perform com Operations or Packs. Because of the large amount of ir propagates events through the Pac to sets of Operations, effectively interested in an event will therefor Packs and Operations can be added different Packs based on their fur – Transformers: Transfor exampl – Metrics: This pack com – Triggers: Triggers are to of 1000 bytec tion of a simp – Printers: Printers are r dynamic call



Trace Analyzer

generated traces and computes appropriate metrics. This component is itself a pipeline of d other components, and so provides flexibility in trace processing as well as easy modification

ganized hierarchically as Packs or Operations. Pack objects are containers for other Packs, nputations. This hierarchical formation is conceptually clean and allows for modular use of

information that is to be processed for each trace file, an event dispatch system that simply ack tree is impractical. Instead, *J preprocesses this tree and generates a mapping from events "flattening" the hierarchy while keeping the relative ordering of Operations. Only Operations ore receive it, thus eliminating the cost of recursively traversing the tree for each event.

ed, removed or redefined. *J's default Pack tree contains several operations grouped into

formers modify the event objects in order to provide services to subsequent operations. For ole, one transformer maps JVMPI identifiers to unique identifiers.

ntains all of the analyses that are used as part of the dynamic metrics project.

used to invoke other specific operations when a particular criterion is met, such as the execution code instructions. Triggers allow metric values to be computed over various intervals by modifica ple counting trigger rather than by having to modify or create new, complex metric operations. responsible for converting the trace file to other formats. For example, one printer generates a l graph based on the recorded method invocations.

Output Processing

In *J, each operation and/or pack is responsible for generating its own output. This is flexible, but not always convenient, and so the default metric pack is capable of aggregating all metric results and generating a single XML file for the entire set of computations

We provide an XSL stylesheet that allows such XML files to be viewed in common web browsers by converting them to static HTML files.

More interestingly, we have implemented a parser for the XML files that is designed to insert the dynamic metric analysis data into a MySQL database. This database is then used to generate dynamic HTML pages on our website using PHP. The website supports arbitrary user-defined database queries. This dynamic website has the advantage of being very easy to update since its contents immediately reflects changes made to the database. This also facilitates the distribution of dynamic metric results to a group of users, and can form the basis of a benchmark knowledge base.



Credits

Bruno Dufour bdufoul@cs.mcgill.ca Clark Verbrugge clump@cs.mcgill.ca Laurie Hendren hendren@cs.mcgill.ca

> Sable Research Group School of Computer Science McGill University

> www.sable.mcgill.ca

This research was supported by NSERC and FQRNT Thanks to Tobias Simon for the dynamic website and databas interface and Marc–André Dufour for the metric icons

Examples of Dynamic Metrics

size.appRun.value: This metric measures the number of bytecode instructions which are executed at least once, or *touched*, during the entire duration of a program's execution. This measurement is not affected by dead code. It is intended to capture the dynamic size of a program, and does not include startup. HelloWorld: 4 Javac: 26267 Compress: 5084

data.floatDensity.value: This metric measures the number of floating-point operations per 1000 executed bytecode instructions. This measurement captures the relative importance of floating-point operations. HelloWorld: 2.0 Javac: 0.1 Mtrt: 308.5

polymorphism.appReceiverArityCalls.bin: This metric measures the percentage of all calls that occur from a call site with one, two and more than two different receiver types. This measurement captures the amount of polymorphism in a program, and does not include startup. HelloWorld: 100%,0%,0% Javac: 73%,15%,12% Jack: 90%,10%,0%

concurrency.lockDensity.value: This metric measures the average number of lock operations (monitorenter)per 1000 executed bytecode instructions. It captures the amount of synchronization found in a program. HelloWorld: 0.19 Compress: 0.0 Jack: 2.13

memory.byteAllocationDensity.value: This metric measures the number of bytes allocated per 1000 executed bytecode instrucitons. It captures the "memory-hungry" aspect of a program. HelloWorld: 1734 Compress: 11 Javac: 132

More metrics descriptions and values can be found at:

http://www.sable.mcgill.ca/metrics/