

# Execution Traces: A New Domain That Requires the Creation of a Standard Metamodel

Luay Alawneh and Abdelwahab Hamou-Lhadj

Electrical & Computer Engineering Department, Concordia University

**Abstract.** Despite the fact dynamic analysis techniques of software systems have been shown to be useful in many software engineering activities such as software maintenance, software performance, testing, etc., there is no standard format for representing run-time information, which hinders interoperability and sharing of data. Runtime information is typically represented in the form of execution traces. Traces can contain different information, and can contain different types of information depending on what is being traced and the purpose of the trace. In this paper, we argue that traces represent vital knowledge about software that needs to be organized and modeled. We support our arguments by discussing the various types of traces used in the literature. We also discuss the challenges when dealing with execution traces and why a trace metamodel has to be carefully designed to overcome these challenges. We also discuss existing attempts to model execution traces. Finally, we discuss how the Knowledge Discovery Metamodel can be extended to support efficiently the modeling of large and complex execution traces.

**Keywords:** Execution Traces, Trace Metamodel, Knowledge Discovery Metamodel.

## 1 Introduction

An important issue in application modernization is the time it takes to understand how the application is built and why it is built this way. In an ideal situation, any change made to an existing software system must be based on information kept in up-to-date documentation. However, for a variety of reasons, it has been shown, in practice, that maintaining sufficiently good documentation is impractical in many organizations, which renders program comprehension a difficult and tedious task. Reverse engineering techniques aim at reducing the impact of this problem by recovering high-level views of the system from low-level implementation details. These views can be used by software engineers to understand the main aspects of the system before diving into the details.

Reverse engineering tools can be grouped into two categories depending on whether they focus on static analysis of the system or on the understanding of its dynamic characteristics. Static analysis techniques operate on the source code to extract a system's main components and their relations. Dynamic analysis, which is the focus of this paper, focuses on the analysis of the behavioural aspects of a system through the analysis of run-time information.

Run-time information is typically represented using execution traces. There exist, however, many types of traces that vary in their structure, the contained information, and the level of abstraction of the contained information. An execution trace can be used to describe the interacting modules involved in a particular scenario, or may be detailed to capture the performed statements in each module's procedure. Examples of such traces include routine-call traces, statement-level traces, traces of inter-process communication, etc.

Although this paper focuses on reverse engineering, dynamic analysis has been shown to be useful in many software engineering activities such as program comprehension, runtime monitoring and performance analysis, testing, fault detection and intrusion detection, etc. There exist many tools for execution trace visualization and analysis. However, these tools use different format for representing traces, which hinders interoperability. This is attributable to a lack of a standardized way for representing execution traces despite the increasing attention to dynamic analysis techniques of software in recent years.

In this paper, we argue that execution traces form a new domain knowledge that needs to be organized and modeled. We discuss the most important types of execution traces, their applications, and their structures. We also discuss how these approaches deal with the trace size problem. Finally, we discuss existing metamodels such as the Knowledge Discovery Metamodel [1] from OMG [2], the UML metamodel [3], etc., and their limitations.

## 2 Execution Traces as a New Domain

Execution traces represent the sequence of execution in a running software system at different levels of abstraction. Traces may exist in different structures according to the degree of abstraction requested by the program analyzer. For example, statement-level traces are linear and contain single executed statements. Routine-level traces depict the sequence of routine calls for a program run and are often represented as a tree structure. Inter-process execution traces capture the interactions between different processes in terms of message passing, and they can be modeled using a graph.

In the following, we present an overview of different types of execution traces, their structure and their application.

### 2.1 Statement-Level Traces

Statement-level traces contain the executed statements of a program run according to a coverage that is usually specified by the user. This type of traces can profile a complete behaviour of the software system and can be used to extract information regarding the flow of control during execution and the dependencies among the executed statements. Also, it helps in maintenance activities such as bug fixing by identifying the cause of the problem.

One of the main challenges when using statement-level traces is the sheer size of the generated trace which can reach millions of events. Therefore, a metamodel that represents this type of traces must be built with scalability in mind. One simple approach to achieve this is to represent the repetitive information contained in a trace only once.

Zhang et al. present a more advanced trace compaction technique called Whole Execution Traces (WET) [4]. WET represents a compressed whole execution trace which captures complete profile information that covers control flow, variable values, variable memory addresses, and control and data dependencies. The framework can respond to a wide range of queries that may require single or multiple types of profile information in a fast and easy manner.

WET is a labelled graph that assigns a unique timestamp to every single instance for each executed statement in order to track the ordering of execution. Furthermore, the paper presents compression techniques that, according to the authors, reduce the size of WET effectively. Moreover, the paper shows some benchmark values that prove the efficiency of the used compression techniques.

## 2.2 Routine Call Traces

Routine-call traces and their object-oriented counterpart “method-call traces” capture the sequence of routine calls in an execution path. This type of traces is at a higher level of abstraction compared to statement-level traces and can be used to reveal significant information about the system’s different scenarios. Routine-call traces, which are represented as tree structures, can be used in many maintenance activities such as bug-fixing [5], and feature-location [5]. Moreover, they can provide useful information that can enable software engineers to perform many activities in an efficient manner including restructuring, refactoring, and code optimization. Like any other type of execution trace, routine-call traces can grow dramatically in size and complexity. Therefore, compression and reduction techniques should be applied in order to utilize them effectively. In the following we present several research studies related to routine and method call traces.

Taniguchi et al. [7] proposed a new method for reverse engineering of UML sequence diagrams based on method-call execution traces. The purpose of this technique is to facilitate the understanding of object-oriented systems. Furthermore, the paper presents a set of four compaction rules in order to solve the problem of large execution traces. The proposed rules can be summarized as follows:

1. Rule 1 finds identical method-call subtrees and represents them in one subtree. Identical method-call subtrees share the same objects. By using this rule, the original call tree can be reconstructed.
2. Rule 2 finds repetitive method calls structures that correspond to different objects and then compacts them by unifying the objects. Although this method can detect repetitive structures that cannot be captured by the first rule, the original subtrees cannot be reconstructed using this rule.
3. Rule 3 detects similar subtrees that may not have the same exact method calls, i.e. one subtree may include one or more method calls than the other. In this case, the compaction is done by using the subtree that contains more method calls.
4. Rule 4 targets the compaction of recursive method calls by combining different recursive calls into one node.

In [8], Wang et al. proposed a new approach for threat model-driven security testing in order to detect threats at runtime. The approach uses UML sequence diagrams in

order to model threats to security policies. A security policy defines what actions should be permitted or refused by the system. In the design phase, threat scenarios are constructed using sequence diagrams. These scenarios depict any expected threat to the system. In order to reduce the size of the trace, the source code is only instrumented with essential information relevant to threat scenarios. Finally, these constructed scenarios are used to verify if any execution trace matches any of the threat scenarios.

Salah et al presented a study [9] that targets program comprehension. It presents a hierarchy of dynamic views composed of different tools for program execution trace analysis. The hierarchy includes feature-interaction, feature-implementation, class-interaction, and object-interaction views. In the same field of program comprehension, Apiwattanapong et al. [10] proposed a new approach for impact analysis of software changes based on dynamic analysis. The presented technique uses only essential dynamic information collected from method-call execution traces.

Finally, Hamou-Lhadj and Lethbridge [11] presented a technique for large execution trace summarization that can be applied to enable top-down analysis of traces as well as the recovery of the behavioural design model of the system. Additionally, the paper proposes a new metric for detecting utility methods which are considered as implementation details that can be removed in the process of abstracting out the main content for large traces.

### 2.3 Inter-process Level Traces

Inter-process communication execution traces capture the interactions among different processes in a software system. Processes may reside on the same computer or different computers. Also, this type of traces captures the communication among threads living within the same process. The main challenge when analyzing this type of execution traces, in addition to size and complexity, is that different executions for the same scenario could generate different traces, which makes it difficult to study this type of traces. The variation in the generated execution traces is due to the non-deterministic behaviour of multi-threaded applications.

Moe et al. [12] apply dynamic analysis through runtime execution traces in order to understand the behaviour of distributed software systems. They propose a method to support the understanding of distributed systems based on the analysis of execution traces at the remote procedure calls level. Also, the work provides a tool for the visualization of the processed execution traces. According to the authors, this work, when applied during the maintenance phase, can help in detecting design flaws, configuration and performance problems.

Bensalem et al. [13] presented an algorithm that uses a single execution trace of multithreaded programs in order to detect occurrences of deadlocks. An advantage of the proposed algorithm is the ability to detect deadlocks in running programs even when examining a deadlock-free execution trace.

### 2.4 System Call Level Traces

A system-call is a request to the kernel by a user-level program in order to be permitted to perform a set of predefined operations that the requesting program does not

possess the required permissions to execute on its own. A system call trace is the sequence of calls made to the system by a running process.

System-call traces can be used to detect and control programs by verifying that each system call conforms to a policy that confirms a program's normal behaviour. Many research works on intrusion detection such as [14-16] use system-call level traces in order to detect and determine anomaly behaviour.

Another application for system-call traces is performance monitoring. Burns et al. [17] used system call execution traces to extract the logical block addresses of a file which are generated over a long period of time in order to evaluate the file system performance.

### 2.5 Execution Traces for Performance Analysis

The increasing size and complexity of software systems require planning for better memory management and CPU processing times. This led to the development of software analysis tools, usually known as profilers, which help in pinpointing execution bottlenecks and aid code optimization consequently. Moreover, the analysis provided by these tools can benefit in decreasing the execution time and reducing the resource utilization such as physical and virtual memory. The main weakness of this type of tools is the overhead introduced from the statically instrumented executed statements. However, Dynamic instrumentation [18] can be used in order to obtain a very low profiling overhead.

Harkema et al. [19] presented a Java Performance Monitoring Toolkit (JPMT) for analyzing the performance of Java applications. It uses event traces such as thread creation, method invocations and locking contention which are annotated by performance attributes such as timestamps in case of method invocations. Instrumentation overhead is overcome by only instrumenting for the types of events requested by the user. Additionally, JPMT supports visualization of event traces and provides the ability of querying for certain types of events.

In his work [20], Putrycz presents a novel approach for analyzing performance in COTS-based systems which uses low-level trace analysis in order to understand the interactions between the communicating components. Pahl et al [21] presents a service-specific approach for performance evaluation of model-driven developed services. This work presents a new approach for instrumentation of model-based languages in order to collect performance-relevant time information at execution time from specific model elements such as services and flow operators.

## 3 Existing Metamodels

There exist several metamodels that are used to capture runtime execution traces such as Compact Trace Format (CTF) [28] and UML [3]. UML sequence diagrams can be used to capture procedure calls among different objects. The problem with the existing metamodels is their inability to model all types of execution traces captured from different software architectures and the lack of support to trace compaction techniques. In this section, we present some of the existing metamodels that are being used or can be used to model different types of execution traces.

### 3.1 Compact Trace Format

Hamou-Lhadj et al. [28] developed a metamodel called the Compact Trace Format (CTF) to model traces of routine (method) calls. CTF was designed to deal with the enormous size of typical traces based on the idea that dynamic call trees can be turned into ordered directed acyclic graphs, where repeated sub-trees are factored out. CTF supports traces defined at different levels of abstraction including object, class and package level. It also supports the specification of threads of execution. Additional information such as timestamps and routine execution time are added to enable profilers to use CTF.

Trace data conforming to CTF can be expressed using GXL [29] or any other data 'carrier' language. However, the authors suggest using a compact representation in order to support the compactness objective of CTF. CTF is lossless such that the original trace can be reconstructed from its compact form.

### 3.2 Unified Modeling Language

UML is a modeling language adopted by OMG in 1997 that enables software designers to specify, visualize, and document software models. These models are abstract representations of the implementation details of software systems. The UML metamodel is based on the Meta Object Facility [30] (MOF) language. MOF defines an abstract language and framework for specifying, constructing and managing technology neutral metamodels.

UML diagrams are classified into two categories: structural and behavioural diagrams. The latter includes a subset known as interaction diagrams. The structural diagrams include those that capture the static structure of software systems such as class and package diagrams. The class and package diagrams help in building metamodels that capture execution traces. On the other hand, behavioural diagrams depict the dynamic behaviour of software systems. Behavioural diagrams include use case diagram, activity diagram, state machine diagram, sequence diagram and others.

The sequence diagram shows object interactions arranged in a time sequence. Sequence diagrams identify the communication required to fulfill an interaction. Moreover, they show the objects that participate in an interaction and the messages used to trigger the interactions among the objects.

There exist some research works that used UML sequence diagram to model runtime execution traces. Briand et al [31] proposed a framework for reverse engineering of UML sequence diagrams using execution traces. This work defines a metamodel for execution traces and maps the execution trace elements to its corresponding sequence diagram elements. The work uses code instrumentation to probe the parts of code that will be used to generate the execution trace. In [32], Delamare et al. used UML 2.0 sequence diagrams to capture the program state from its execution traces for the purpose of program understanding.

In [33], the authors used UML State Machine diagrams as the basis for their approach to runtime verification of Java programs. The approach studies the temporal order of message receiving based on consistency checking between the behaviour of state machine diagrams and the program execution traces.

### 3.3 Knowledge Discovery Metamodel

The Knowledge Discovery Metamodel (KDM) [1] is a metamodel that targets a wide-spread set of software applications, platforms and programming languages such as modern enterprise applications which involve multiple technologies and programming languages. The goal of KDM is to facilitate the integration between different tools that capture information about complex enterprise applications. The structure of KDM offers a common interchange format, using XMI schema, which allows interoperability between existing tools and their models. Moreover, KDM captures the physical and logical software assets at various levels of abstraction as entities and relations. This nominates it as a favorable basis for different software domains.

KDM is designed based on the separation of concerns principle in order to enable different compliant tools to support the same or compatible metamodel subsets. This modular structure of the metamodel allows a tool vendor to select only its desired or needed parts of the metamodel. Furthermore, the structure of KDM consists of different packages that represent each domain in enterprise applications. This modular structure allows for the extensibility of the KDM metamodel by adding new domains to the metamodel as needed.

This structure of KDM means that users need only to learn about the domain of their interest. For example, the Structural domain provides users with information about the architectural elements from the source code of the target system. On the other hand, the Business Rules domain provides users with behavioural elements of the system such as features or process rules.

The KDM metamodel is organized in four different layers. The KDM infrastructure layer defines the basis for the KDM metamodel. Its packages are used by the packages in the other layers. The Program Elements Layer defines a large set of metamodel elements whose purpose is to provide a language-independent intermediate representation for various constructs determined by common programming languages. The Runtime Resource Layer describes common patterns for representing the operating environment of existing software systems. Finally, the Abstraction Layer defines a set of metamodel elements whose purpose is to represent domain-specific and application specific abstractions, as well as the engineering view of the existing software system.

## 4 Proposed Execution Trace Metamodel

Runtime execution traces represent a separate domain in software modernization. They provide proper understanding of the different parts of the system under study. Also, they can facilitate different software maintenance and performance monitoring activities. Execution traces may exist in different levels of abstraction. The objective of this work is to support execution traces in all levels of abstraction and to define a standardized form for execution traces that supports meaningfulness, abstraction and expressiveness.

Execution traces can be generated using a technique known as program instrumentation. Instrumentation of the source code should be performed properly in order to generate an execution trace, at a certain level of abstraction, which can be applied feasibly in order to achieve the goal of the analysis task.

The proposed metamodel should be flexible to cover the aforementioned types of execution traces. Therefore, it should be based on a metamodel that supports extensibility in order to cope with newer types of traces. Our discussion on KDM shows that it can be a proper candidate for our proposed metamodel because of the following advantages:

1. KDM is a metamodel that targets a widespread set of software applications, platforms and programming languages such as modern enterprise applications which involve multiple technologies and programming languages.
2. Separation of concerns concept. This helps in extending KDM to support different domains by adding new packages to the metamodel.
3. KDM uses the XMI schema to store the software artifacts. XMI is an OMG standard for exchanging metadata information via Extensible Markup Language (XML). It can be used for any metadata whose metamodel can be expressed in Meta-Object Facility (MOF) such as UML.
4. KDM captures the physical and logical software assets at various levels of abstraction as entities and relations.
5. KDM metamodel defines program element entities and their relationships which can play a main role in building a comprehensive execution trace metamodel. Executed traces can be mapped easily to their corresponding program elements since KDM assigns a unique identifier for each program element.

We are interested in the KDM Runtime Resource Layer because it represents the dynamic structures, instances of logical entities and their relationships, which exist at runtime such as processes and threads. Therefore, a new package to represent the Runtime Execution traces can be created in this layer. Figure 1 depicts the structure of KDM packages along with our new Trace package that will represent the execution traces domain.

The advantages of our approach are manifold and can be summarized as:

1. Our metamodel will utilize the structure of KDM. Therefore, runtime execution traces can be exchanged easily among different analysis tools.
2. The new Trace package will reuse various KDM packages such as Core, Code and Action.
3. The execution trace model can follow the Directed Acyclic Graph structures. Therefore, different graph reduction and summarization techniques can be applied to our metamodel.
4. Polymorphism and dynamic binding in object oriented systems will be supported in our metamodel easily since KDM assigns a unique identifier to every element in the source code. Therefore, each method will be instrumented with its KDM unique identifier. Thus, a method call in the execution trace can be linked to its class using its unique identifier.
5. Processes and Threads are already supported in KDM and will be reused in our Trace package.
6. The Trace package can be extended to support newer types of execution traces easily due to the extensibility nature of KDM.
7. The new metamodel can be integrated easily with several visualizations schema such as GXL.



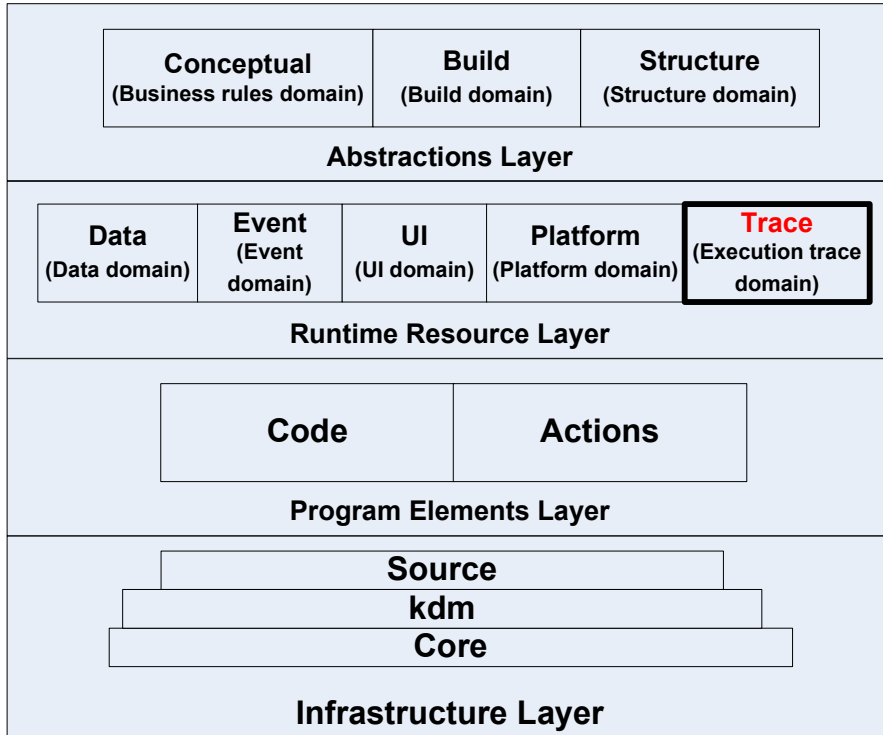


Fig. 1. Updated KDM Structure with Trace Package

## 5 Conclusion and Future Work

This paper presented runtime information through execution traces as a new domain in software engineering supported by several research studies that target or utilize execution traces to achieve their objectives. We discussed a few metamodels that are used to capture execution traces. Our discussion showed that the available metamodels lack the possibility of capturing all types of execution traces. Moreover, these metamodels except for [28] do not apply trace compaction techniques. Finally, we proposed building a new metamodel based on KDM for its numerous advantages. The resulting metamodel should be able to model any type of execution traces in a compact form.

Our future work will focus on building the new metamodel for the execution trace domain. We will continue studying all the available types of execution traces in order to support them in our metamodel.

## References

1. Object Management Group. Knowledge Discovery Metamodel: KDM Version 1.1 Beta 3 (March 2008)
2. OMG: Object Management Group, <http://www.omg.org/>

3. Object Management Group. Unified Modeling Language: Infrastructure and Superstructure, Version 2.0, formal/2007-11-04 (November 2007)
4. Zhang, X., Gupta, R.: Whole execution traces and their applications. *ACM Transactions on Architecture and Code Optimization (TACO)* 2(3), 301–334 (2005)
5. Cleve, H., Zeller, A.: Locating causes of program failures. In: *ACM/IEEE International Conference on Software Engineering, ICSE (2005)*
6. Liu, D., Marcus, A., Poshyvanyk, D., Rajlich, V.: Feature Location via Information Retrieval based Filtering of a Single Scenario Execution Trace. In: *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering*, pp. 234–243 (2007)
7. Taniguchi, K., Ishio, T., Kamiya, T., Kusumoto, S., Inoue, K.: Extracting Sequence Diagram from Execution Trace of Java Program. In: *Proceedings of the Eighth International Workshop on Principles of Software Evolution*, pp. 148–154 (2005)
8. Wang, L., Wong, E., Xu, D.: A Threat Model Driven Approach for Security Testing. In: *Proceedings of the Third International Workshop on Software Engineering for Secure Systems*, pp. 111–112 (2007)
9. Salah, M., Mancoridis, S.: A Hierarchy of Dynamic Software Views: From Object-Interactions to Feature-Interactions. In: *Proceedings of the 20th IEEE International Conference on Software Maintenance*, pp. 72–81 (2004)
10. Apiwattanapong, T., Orso, A., Harrold, M.: Efficient and precise dynamic impact analysis using execute-after sequences. In: *Proceedings of the 27th international conference on Software engineering (2005)*
11. Hamou-Lhadj, A., Lethbridge, T.: Summarizing the Content of Large Traces to Facilitate the Understanding of the Behaviour of a Software System. In: *Proceedings of the 14th IEEE International Conference on Program Comprehension*, pp. 181–190 (2006)
12. Moe, J., Carr, D.: Understanding Distributed Systems via Execution Trace Data. In: *Proceedings of the 9th International Workshop on Program Comprehension*, pp. 60–67 (2001)
13. Bensalem, S., Havelund, K.: Scalable deadlock analysis of multi-threaded programs. In: *Proceedings of the Parallel and Distributed Systems: Testing and Debugging (PADTAD) Track of the 2005 IBM Verification Conference*. Springer, Heidelberg (2005)
14. Varghese, S.M., Jacob, K.P.: Anomaly Detection Using System Call Sequence Sets. *Journal of Software* 2(6) (2007)
15. Fetzer, C., Suesskraut, M.: SwitchBlade: Enforcing Dynamic Personalized System Call Models. In: *Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems (2008)*
16. Goel, A., Feng, W., Maier, D.: Automatic high-performance reconstruction and recovery. *The International Journal of Computer and Telecommunications Networking* 51(5), 1361–1377 (2007)
17. Burns, R., Rees, R., Peterson, Z., Darrell, D.E.: Allocation and Data Placement Using Virtual Contiguity, iNIST/SSRC/01-001, pp. 1–6 (2001)
18. Cantrill, B.M., Shapiro, M.W., Leventhal, A.H.: Dynamic instrumentation of production systems. In: *Proceedings of the USENIX Annual Technical Conference 2004 on USENIX Annual Technical Conference (2004)*
19. Harkema, M., Quartel, D., van der Mei, R., Gijssen, B.: JPMT: a Java performance monitoring tool. In: *Proceedings of the 13th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation (2003)*
20. Putrycz, E.: Using trace analysis for improving performance in COTS systems. In: *Proceedings of the 2004 conference of the Centre for Advanced Studies on Collaborative research*, pp. 68–80 (2004)
21. Pahl, C., Boskovic, M., Hasselbring, W.: Model-Driven Performance Evaluation for Service Engineering. In: *Proceedings of the 2nd European Conference on Web Services Workshop on Web Services Technology (2007)*

22. McGavin, M., Wright, T., Marshall, S.: Visualisations of Execution Traces (VET): An Interactive Plugin-Based Visualisation Tool. In: Proceeding of the 7th Australasian User Interface Conference, pp. 153–160 (2006)
23. Fischer, M., Oberleitner, J., Gall, H., Gschwind, T.: System Evolution Tracking through Execution Trace Analysis. In: Proceedings of the 13th International Workshop on Program Comprehension, pp. 237–246 (2005)
24. Cornelissen, B., Zaidman, A., Holten, D., Moonen, L., van Deursen, A., van Wijk, J.: Execution trace analysis through massive sequence and circular bundle views. *Journal of Systems and Software* 8(12) (2008)
25. de Kergommeaux, J.C., de Oliveira Stein, B.: Pajé: An Extensible Environment for Visualizing Multi-threaded Programs Executions. In: Proceedings of the 6th International EuroPar Conference on Parallel Processing, pp. 133–140 (2000)
26. Roberts, J., Zilles, C.: TraceVis: An Execution Trace Visualization Tool. In: Proceedings of the Workshop on Modeling, Benchmarking, and Simulation (2005)
27. Malnati, G., Cuva, C.M., Barberis, C.: JThreadSpy: teaching multithreading programming by analyzing execution traces. In: Proceedings of the Parallel And Distributed Systems: Testing and Debugging Conference (2007)
28. Hamou-Lhadj, A., Lethbridge, T.C.: A Metamodel for Dynamic Information Generated from Object-Oriented Systems. *Electronic Notes Theoretical Computer Science*, vol. 94, pp. 59–69. Elsevier Press, Amsterdam (2004)
29. Winter, A., Kullbach, B., Riediger, V.: An Overview of the GXL Graph Exchange Language, Revised Lectures on Software Visualization. In: International Seminar, pp. 324–336 (2001)
30. Object Management Group. Meta Object Facility (MOF) Specification (2000)
31. Briand, L.C., Labiche, Y., Miao, Y.: Towards the Reverse Engineering of UML Sequence Diagrams. In: Proceedings of the 10th Working Conference on Reverse Engineering (2003)
32. Delamare, R., Baudry, B., Traon, Y.L.: Reverse-engineering of UML 2.0 Sequence Diagrams from Execution Traces. In: Workshop on Object-Oriented Reengineering at ECOOP (2006)
33. Li, X., Qiu, X., Wang, L., Lei, B., Wong, W.E.: UML State Machine Diagram Driven Runtime Verification of Java Programs for Message Interaction Consistency. In: Proceedings of the 23rd Annual ACM Symposium on Applied Computing (ACM SAC 2008), pp. 384–389 (2008)