

Towards a Scalable Representation of Run-time Information: The Challenge and Proposed Solution

Abdelwahab Hamou-Lhadj

Electrical and Computer Engineering, Concordia University, 1455 de Maisonneuve West
Montreal, Canada

Abstract. An important issue in application modernization is the time and effort needed to understand existing applications. Understanding the dynamic aspects of software is a difficult task, further complicated by the lack of a scalable way of representing the extracted knowledge. The behavior of a software system is typically represented in the form of execution traces. Traces, however, can be extraordinary large. Existing metamodels such as the Knowledge Discovery Metamodel and the UML metamodel provide limited support for handling large execution traces. In this paper, we describe a metamodel called the Compact Trace Format (CTF) for efficient modeling of traces of routine (method) calls generated from multi-threaded systems. CTF is intended to facilitate the interoperability among modernization tools that focus on the analysis of the behavior of software systems. CTF is designed to be easily extensible to support other types of traces.

1 Introduction

Software modernization is defined as the process of understanding and evolving existing software systems for the purpose of improving their various facets [1].

Understanding a software system requires both static and dynamic analysis techniques. The former focuses on exploring the structure of the system by analyzing its source code, whereas the latter, which is the focus of this paper, provides insight into the system's behavioral properties. Both approaches aim at extracting knowledge about the system under study that can later be fed to a modernization tool for further analysis.

There exist several standards for representing the knowledge extracted from software systems, among which the most recent is the Knowledge Discovery Metamodel (KDM) [10] supported by the Object Management Group (OMG). Others include the UML 2 metamodel [17], the Dagstuhl Middle Metamodel (DMM) [12], etc.

However, while existing metamodels provide a full range of constructs for the modeling of the static aspects of the system (i.e. its static components and the way they interact), they do not support efficient representation of the system's behavioral features, typically represented in the form of execution traces. Traces have historically been difficult to work with since they may contain millions of events. The challenge is to develop a trace format that scales up to handle large and complex traces.

In this paper, we describe ongoing work towards the development of a scalable format for representing execution traces. A particular format called CTF (Compact Trace Format) is proposed. CTF focuses on representing traces of routine (methods) calls generated from multi-threaded systems. It is built with scalability in mind using graph theory concepts. In addition, CTF metamodel is defined in such a way that it can be easily extended to support other types of traces and constructs.

The remaining of this paper is organized as follows: In Section 2, we discuss the requirements that guided us in the development of CTF. In Section 3, we present the CTF metamodel and its syntactic form. CTF semantics are presented in the appendix section.

2 Requirements for a Trace Exchange Format

Requirements for an effective exchange format have been the subject of many studies [2], [15], [18]. In this section, we present the ones that a trace exchange format should fulfill in order to facilitate its adoption. These requirements were used to guide the development of CTF.

2.1 Scalability

The adoption of a trace exchange format will greatly depend on its ability to support large-sized traces. Trace, however, tend to contain many repetitions due to the presence of loops and recursions in the source code. They also contain non-contiguous repetitions known as trace patterns. The scalability of a trace exchange format can, therefore, be significantly improved if repeated events are factored out and represented only once. For example, the trace of routine calls described in Figure 1a can be transformed into an ordered directed acyclic graph (DAG) represented in Figure 1b by representing the repeated subtree rooted at “B” only once. This technique was first introduced by Downey et al. in [3] to improve tools that manipulate trees. It was also used by Larus [11] and Reiss and Renieris [14] to encode traces with the objective of saving disk space.

It should be noticed, however, that the graph must be *ordered* in order to be able to restore the initial order of calls. In addition, the graph representation of a trace does not necessarily result in a loss of information associated with individual nodes of the tree such as timing information commonly collected when generating a trace. The simplest solution is to augment the nodes of the graph with ordered collections that holds the information describing individual nodes of the tree.

In previous work [6], we applied this compaction technique to over thirty execution traces and showed that it can reach a 97% compression ratio (i.e. the graph contains only 3% of the number of nodes of the tree). This has led us, as we will see in Section 3, to build CTF using the ordered DAG as its main mechanism.

2.2 Completeness

This requirement consists of having an exchange format include all the necessary information needed during the exchange: The data to exchange as well the metamodel that describes the structure of the data. The rationale behind this is to enable tools to check the validity of the instance data against the metamodel. To address this requirement, we need to select a syntactic form language (i.e. the language that “carries” CTF instance data) that is designed to support the exchange of the instance data as well as the metamodel. In Section 3.2, we discuss possible syntactic forms that could be used with CTF.

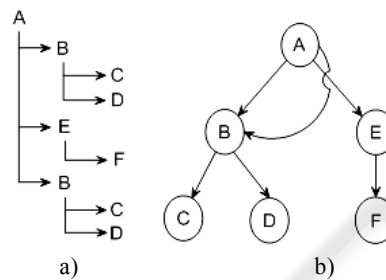


Fig.1. a) An example of a call tree. b) The compact form of the call tree.

2.3 Extensibility

The data model of a trace exchange format must be flexible enough to support new types of traces, language specific entities and properties, and some properties that are specific to the trace analysis tool that uses the format. CTF addresses this requirement by adopting an open design based on abstract classes that allow new constructs to be easily added.

2.4 Tool Support

In order to facilitate the adoption of an exchange format by tool builders, we need to develop well-defined mechanisms that facilitate the manipulation of traces. First, we need to design procedures that ensure that the information exchanged is represented without any alteration. Second, we need to develop algorithms for on the fly generation of traces in the new format. Finally, we need to create converters that will convert other commonly used formats into the new format to facilitate the transition to the new exchange format. Since CTF is still an ongoing project, this requirement will be addressed in future work.

maintains a collection of timestamps and another one that stores the execution times of the individual routine calls represented by this node.

Edges (i.e. calls) are depicted using the TraceEdge class. An edge links two nodes of the DAG. Nodes that are repeated contiguously due to the presence of loops or recursion in the source code are collapsed into one single node; the attribute “repetitions” of the class TraceEdge records the number of repetitions.

As its name indicates, the class TracePattern represents patterns of execution invoked in the trace. A trace pattern is defined as a sequence of events that is repeated non-contiguously in the trace. They are used by software engineers as a means to uncover key domain concepts from the trace. The common hypothesis is that the sequence of calls that appears in various places in the trace might encapsulate some knowledge about the system such as a particular aspect of an algorithm, etc. [9], [16]. Software engineers often associate a textual description to trace patterns that are deemed important. The class TracePattern uses an attribute called “description” to capture this textual description.

A node can either be a routine call node (RoutineCallNode), a method node (MethodCallNode), or a control node (ControlNode). A RoutineCallNode object represents a call to a routine that is not a method of a class, whereas a MethodCallNode stands for a method invocation. Method calls are considered as routine calls except that they may contain additional information such as the objects (represented using the class Object) on which the methods are invoked.

The above metamodel relies on a string label to identify trace events (i.e. the routines invoked). It does not provide any other static information about the routines (or methods) invoked such as the parameter list and return type. The reason is that CTF is intended to work with existing metamodels such the UML metamodel, KDM, or DMM. For example, DMM defines two classes called Routine and Method that describe the static elements of a routine and a method respectively. Assuming that CTF is used with DMM, the classes RoutineCallNode and MethodCallNode will need to be linked to DMM classes Routine and Method in order to retrieve static information. This linkage is not defined in this paper and will be the subject of future improvements of CTF. The advantage of such design is that it allows CTF to be adopted by static analysis tools that use the aforementioned metamodels.

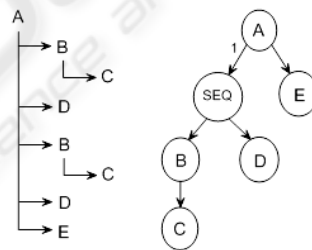


Fig.3. The control node SEQ is used to represent the contiguous repetitions of the subtrees rooted at B and D.

Control nodes represent extra information that might be used by software engineers to customize parts of the trace. In particular, we define two control nodes: SequenceNode and RecursionNode. A SequenceNode object is used to represent

contiguous repetitions of multiple sequences of events. Figure 3 illustrates how such node is used to avoid repeating the sequence “BC and D” twice. The label “SEQ” is used in CTF to identify a SequenceNode object.

A recursive sequence of calls is represented using another control node called RecursionNode, which in turn refers to the recursively repeated sequence of calls. Figure 4 shows a recursion occurrence node labeled “REC” that is used in CTF to collapse the recursive repetitions of the node B. It should be noted that the number of repetitions is captured through the attribute “repetitions” of the class TraceEdge since an edge will be created between the routine “A” and the recursive sequence.

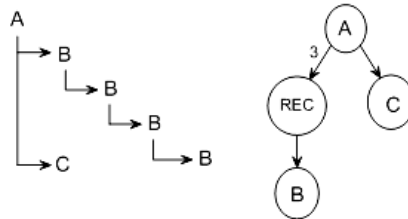


Fig.4. The control node REC is used in CTF to represent the recursive repetitions of a sequence of calls.

The class Thread represents the thread that executes the corresponding portion of the trace (represented by the class Node). Threads are identified using unique thread names since this is a common practice in languages such as Java and C++. We do not distinguish between thread start/end routines from other routines for simplicity reasons. An obvious extension to this model is to fully support various multi-thread communication mechanisms.

3.2 Syntactic Form

CTF instance data can be carried using a syntactic form that supports the representation of graph structures. There exist several languages that satisfy this requirement, which differ essentially on whether they rely on XML or not. In this paper, we discuss how the Graph Exchange Language (GXL) [8] and the Tuple Attribute (TA) [7] can be used with CTF. The choice of these languages is due to the fact that they are widely used by the reverse engineering research community. In addition, both languages support the exchange of the instance data as well as the metamodel; this is compliant with the completeness requirement discussed in Section 2.2.

A GXL file consists of XML elements for describing nodes, edges, attributes, etc. It was designed to supersede a number of pre-existing graph formats such as GraX [4], TA [7], and RSF [13]. GXL has been widely adopted as a standard exchange format for various types of graphs by both industry and academia.

However, an XML-based representation of a trace would tend to be much larger than necessary due to the use of XML tags and the explicit need to express the data as XML nodes and edges. The compactness benefits of a trace exchange format discussed earlier would therefore be partially cancelled out by representing it using XML. One reasonable alternative to GXL is the Tuple Attribute (TA) [7], which

would substantially reduce the space required by a CTF trace (since it is not based on XML). The TA language was originally developed to help visualize information about software systems. It has been used as a model interchange format in several contexts and has a reasonable tool support despite the fact that it is not XML-based.

We are still in the process of testing CTF in order to determine the most suitable syntactic form that represents its instance data efficiently.

4 Conclusions

Application modernization requires the representation of knowledge about the system under study. Existing metamodels such as KDM, UML, and DMM are designed mainly to model the structure of a software system. They lack an efficient representation of its behavioral properties. In this paper, we presented CTF (Compact Trace Format), an exchange format for representing traces of routine (method) calls. To deal with the vast size of typical traces, we designed CTF based on the idea that dynamic call trees can be turned into ordered directed acyclic graphs, where repeated subtrees are factored out. CTF, as described in this paper, is a metamodel. Trace data conforming to CTF can be expressed using GXL, TA, or any other data “carrier” language. However, we suggest using a compact representation since doing otherwise would somewhat defeat the compactness objective of CTF.

While CTF covers a significant gap in terms of exchanging traces of routine calls, dynamic analysis is a highly versatile process that has a large number of needs including needs for dynamic information that is not necessarily supported by CTF. Therefore, the main future work is to work towards enhancing CTF in order to support other types of traces and constructs as well as testing CTF using large traces.

References

1. Bézivin J., 2004. Model Engineering for Software Modernization. In *WCRE'04, 11th Working Conference on Reverse Engineering*.
2. Bowman T., Godfrey M. W., and Holt R. C., 1999. Connecting Architecture Reconstruction Frameworks. In *CoSET'99, 1st International Symposium on Constructing Software Engineering Tools*.
3. Downey J. P., Sethi R., Tarjan R. E., 1980. Variations on the Common Subexpression Problem. *Journal of the ACM*. 27(4).
4. Ebert J., Kullbach B., Winter A., 1999. GraX – An Interchange Format for Reengineering Tools. In *WCRE'99, 6th Working Conference on Reverse Engineering*.
5. Hamou-Lhadj A. and Lethbridge T., 2003. The Compact Trace Format, In *ATEM'03, 1st International Workshop on Metamodels and Schemas for Reverse Engineering*.
6. Hamou-Lhadj A. and Lethbridge T., 2005. Measuring Various Properties of Execution Traces to Help Build Better Trace Analysis Tools. In *ICECCS'05, 10th IEEE International Conference on Engineering of Complex Computer Systems*.
7. Holt R. C., 1998. An Introduction to TA: The Tuple Attribute Language. *Department of Computer Science, University of Waterloo and University of Toronto*.
8. Holt R. C., Winter A., Schürr A., 2000. GXL: Toward a Standard Exchange Format. In *WCRE'00, 7th Working Conference on Reverse Engineering*.

9. Jerding D., Rugaber S., 1997. Using Visualisation for Architecture Localization and Extraction. In *WCRE'97, 4th Working Conference on Reverse Engineering*.
10. Knowledge Discovery Metamodel (KDM) 1.0 Specifications, 2007: <http://www.omg.org/cgi-bin/doc?ptc/2007-03-15>
11. Larus J. R., 1999. Whole program paths. In *PLDI'99, Conference on Programming language design and implementation*.
12. Lethbridge T. C., 2003. The Dagstuhl Middle Model: An Overview. In *ATEM'03, 1st International Workshop on Metamodels and Schemas for Reverse Engineering*.
13. Müller H. A., Klashinsky K., 1988. Rigi – A System for Programming in-the-large. In *ICSE'88, 10th International Conference on Software Engineering*.
14. Reiss S. P. and Renieris M., 2001. Encoding program executions. In *ICSE'01, 23rd International Conference on Software Engineering*.
15. St-Denis G., Schauer R., and Keller R. K., 2000. Selecting a Model Interchange Format: The SPOOL Case Study. In *Proc. of the 33rd Annual Hawaii International Conference on System Sciences*.
16. Systä T., 2000. Understanding the Behaviour of Java Programs. In *Proc. of the 7th Working Conference on Reverse Engineering (WCRE)*.
17. UML 2.0 Specifications, 2005: <http://www.omg.org/technology/documents/formal/uml.htm>
18. Woods S., Carrière S. J., Kazman R., 1999. A semantic foundation for architectural reengineering and interchange. In *ICSM'99, 15th International Conference on Software Maintenance*.

Appendix

In this appendix, we present the detailed semantics of CTF.

Class “Trace”

Semantics: An abstract class representing common information about traces generated from the execution of the system.

Attributes:

- startTime: Time - Specifies the starting time of the generation of the trace.
- endTime: Time - Specifies the ending time of the generation of the trace.

Associations: No associations.

Constraints:

- endTime should be greater than or equal to startTime: `self.endTime >= self.startTime`

Class “CallTree (Subclass of Trace)”

Semantics: An object of CallTree represents a trace of routine calls. A routine is defined as any function whether it is in a class or not. Although the class refers to a tree but, in fact, it will be saved as an ordered DAG.

Attributes: No additional attributes.

Associations:

- root: Node[1] - Specifies the root of the call tree.

Constraints:

- The root of a trace must not have parent node: `self.root.incoming ->isEmpty()`
- The root node cannot be an object of ControlNode subclasses:
`not self.root.oclIsTypeOf(SequenceNode)` and
`not self.root.oclIsTypeOf(RecursionNode)`
- The graph needs to be an ordered directed acyclic graph.

Class “TracePattern”

Semantics: An object of the class TracePattern represents a sequence of calls that is repeated in a non-contiguous manner in the trace.

Attributes:

- description: String - Specifies a textual description that a software engineer can assign to a trace pattern.

Constraints: No additional constraints.

Class “Node”

Semantics: Node is an abstract class that represents the nodes of the directed acyclic graph (i.e. compact form of the call tree).

Attributes:

- label: String - Represents the name of the routine represented by this node.
- timestamps: Time[] - Specifies the timestamps of the routines represented by this node.
- executionTime: double[] - Specifies the execution times of the routines represented by this node.

Associations:

- DAG: CallTree [1] - References the Trace for which this node is the root.
- incoming: TraceEdge [*] - Specifies the TraceEdge objects that represent the incoming edges of this node.
- outgoing: TraceEdge [*] - Specifies the TraceEdge objects that represent the outgoing edges of this node.
- Thread [*] - References the Thread objects that represent the thread in which this node is executed.

Constraints:

- The timestamps of the routine calls represented by this node must be sorted in an ascending manner. This guarantees that the graph maintains the sequential execution of the routines.
- The parent nodes of this node cannot be the same as its child nodes and vice-versa since the graph is acyclic.
`self.incoming->excludesAll(self.outgoing)` and
`self.outgoing ->excludesAll(self.incoming)`

Class “TraceEdge”

Semantics: Objects of the TraceEdge class represent edges of the directed acyclic graph.

Attributes:

- repetitions: int - Specifies an edge label that will be used to represent the number of repetitions due to loops and recursion. Default value is zero, i.e., no repetitions.

Associations:

- child: Node[1] - References the node that represents the child node that is pointed to by this trace edge.
- parent: Node [1] - References the node that represents the parent node from which this edge is an outgoing edge.

Constraints:

- The child and the parent nodes must be different. Recursion is represented using the RecursionNode class (Section 3.1):
`self.child <> self.parent`
- The value of the attribute “repetitions” must be greater than or equal to zero
`self.repetitions >= 0`

Class “Thread”

Semantics: Objects of the Thread class represent the thread of execution invoked in the trace.

Attributes:

- name: String - Specifies the name of the thread.

Associations:

- Node[1..*] - References the nodes that are executed in this thread of execution.

Constraints: No constraints.

Class “RoutineCallNode (Subclass of Node)”

Semantics: Objects of the RoutineCallNode represent the routine calls invoked in the trace.

A routine here should not be confused with a method of a class.

Attributes: No additional attributes.

Associations: No additional associations.

Constraints: No additional constraints.

Class “MethodCallNode (Subclass of Node)”

Semantics: Objects of the MethodCallNode represent the method calls invoked in the trace.

Attributes: No additional attributes.

Associations:

- Object [0..1] - References the object, if known, on which the method is invoked.

Constraints: No additional constraints.

Class “Object”

Semantics: This class represents the objects invoked in the trace. In some traces, information about objects may be present; in others such information (and hence instances of this class) may be absent.

Attributes:

- objectID: String - Specifies the object identifier.

Associations:

- MethodCallNode [1..*] - Specifies the methods invoked on this object.

Constraints: No additional constraints

Class “ControlNode (Subclass of Node)”

Semantics: The ControlNode class is an abstract class that is used to specify additional information that can help better structure the trace.

Attributes: No additional attributes.

Associations: No additional associations.

Constraints:

- A control node cannot be the root of the entire trace: self.incoming ->notEmpty()
- A control node must have children: self.outgoing -> notEmpty()

Class “RecursionNode (Subclass of ControlNode)”

Semantics: An object of the RecursionNode is added to represent graph nodes that are repeated recursively. In this case, this object will be labeled ‘REC’.

Attributes: No additional attributes.

Associations:

- repeatedOccurrence: Node[1] - References the subtree that is repeated recursively.

Constraints: No additional constraints

Class “SequenceNode (Subclass of Control Node)”

Semantics: An object of the SequenceNode class is added to represent multiple nodes that are repeated in a contiguous way. In this case, this object will be labeled ‘SEQ’.

Attributes: No additional attributes.

Associations: No additional associations.

Constraints: No additional constraints.