

THE NATIONAL UNIVERSITY
of SINGAPORE

School of Computing
Lower Kent Ridge Road, Singapore 119260

TRA8/06

***SMArTIC: Specification Mining Architecture
with Trace filtering and Clustering***

David LO and Siau-Cheng Khoo

August 2006

Technical Report

Foreword

This technical report contains a research paper, development or tutorial article, which has been submitted for publication in a journal or for consideration by the commissioning organization. The report represents the ideas of its author, and should not be taken as the official views of the School or the University. Any discussion of the content of the report should be sent to the author, at the address shown on the cover.

JAFFAR, Joxan
Dean of School

SMArTIC: Specification Mining Architecture with Trace filtering and Clustering

David Lo and Siau-Cheng Khoo

Department of Computer Science, National University of Singapore,
{dlo, khoosc}@comp.nus.edu.sg

Abstract. Improper management of software evolution, compounded by imprecise, and changing requirements, along with the “short time to market” requirement, commonly leads to a lack of up-to-date specifications. This can result in software that is characterized by bugs, anomalies and even security threats. Software specification mining is a new technique to address this concern by inferring specifications automatically. In this paper, we propose a novel API specification mining architecture called SMArTIC (Specification Mining Architecture with Trace filtering and Clustering) to improve the accuracy, robustness and scalability of specification miners. This architecture is constructed based on two hypotheses: (1) Erroneous traces should be pruned from the input traces to a miner, and (2) Clustering related traces will localize inaccuracies in learning and reduce over-generalization. Corresponding, SMArTIC comprises four components: an erroneous-trace filtering block, a related-trace clustering block, a learner, and a merger. We show through experiment that the quality of specification mining can be significantly improved using SMArTIC.

1 Introduction

Improper management of software evolution commonly leads to a lack of up-to-date specifications (*cf.* [1]). This situation is further aggravated by imprecise, changing requirements, along with the “short time to market” requirement [2]. Low coherence between specifications and implementations can result in software that are characterized by bugs, anomalies and even security threats. There has been continual effort to develop techniques which aim to infer specifications automatically. In recent years, we have also seen a surge within the software engineering research community to adopt dynamic analysis, machine learning and statistical approaches to address these problems, especially in the area of specification discovery [3–10]. These methods are generally termed *specification miners*. In [11], Fox illuminates the use of machine learning to bridge the gap between high level abstractions expressing software engineering problems and low level program behaviors. He points out that some baseline models can be learned automatically to aid in the characterization and monitoring of systems.

Specification miners can be classified into two groups, depending on how the mined specifications are represented: *automaton-based* [6, 3, 9, 5, 7] and *non-automaton based* [10, 8, 4, 12] specification miners.

The work by Ammons *et al.* is a pioneer in automaton-based specification mining [6]. There, a machine-learning approach is employed to discover program specifications

by analyzing program execution traces. Under the assumption that the program being mined must “reveal strong hints of correct protocols” during its execution, Ammons *et al.* demonstrate that correct specifications can be obtained through their technique. Specifically, their technique focuses on mining of specifications which reflect temporal and data dependency relations of a program through traces of its API-client interaction. The specifications discovered model API-client interaction protocols, which are expressed initially as a probabilistic finite state automaton (PFSA). To reduce the effect of errors in training traces, transitions with a low likelihood of being traversed can later be pruned. After pruning, the probabilities are dropped and an FSA is obtained.

In this paper, we leverage on the work performed by Ammons *et al.* in automaton-based specification mining, and explore the art and science behind the construction of such a miner. Specifically, we devise a novel architectural framework that achieves specification mining through pipelining of four functional components: Error-trace filtering, clustering, learning, and automaton merging. We demonstrate that such a system architecture improves the quality of the mining result for two primary reasons:

1. Early identification and filtering of erroneous program execution traces can improve the quality of specification discovery.
2. Over-generalization which occurs at the learning stage can be mitigated by localization of learning process to groups of related program execution traces.

Contrary to other works done in automaton-based specification mining, we choose *probabilistic FSA* (PFSA) instead of (non-probabilistic) FSA as our learning target. PFSA is more expressive than FSA since it provides details on the probabilities of state transitions. This enables detection of frequently-used interaction patterns (*eg.*, “open (read)* close” pattern in a resource-access protocol) or sub-protocols within a specification, analogous to the idea of hotspots found in program execution [13].

We conduct experiments to support our reasoning. Our experiments aim at deriving the API interaction protocol for a client of the Jakarta Commons Net open-source library [14]. We adopt an objective measurement of the quality of our mining architecture through introduction of *accuracy*, *robustness* and *scalability*.

Accuracy refers to the extent of an inferred specification being representative of the actual specification. *Robustness* refers to its sensitivity to errors present in the input data. *Scalability* determines a specification miner’s ability to infer large specifications.

These measurements extend from the existing set of measurements found in the literature; specifically, the measurement of *accuracy* is supported by three kinds of metrics: precision, recall and co-emission. Measurement of *recall* and *precision* are suggested by Nimmer and Ernst [15]. They have been used as a measurement of soundness and completeness of specification discovery. As our specification miner manipulates specifications in PFSA format, it is necessary to also measure if the mined PFSA generates *the same traces at similar frequencies*, and thus places emphasis on similar sub-protocols. This is commonly measured by a metric known as *co-emission* [16].

The outline of the paper is as follows: In section 2, we provide a brief introduction to issues pertaining to experimentation on our specification miners. Section 3 lays down the hypotheses which drive our construction of SMArTIC, and discusses the detail components in SMArTIC. Section 4 describes our experiments on the Jakarta Commons Net [14] open source library. Section 5 describes more comprehensive experiments

and results using various simulated models. We discuss related work in Section 6 and conclude in Section 7.

2 Preliminary

In order to assess the quality of SMArTIC, we need to have an objective means of measuring an output automaton against an expected automaton. To this end, we devise the following simulation model to facilitate experimentation with a variety of automata: We begin with a known API-interaction specification of a software component. This specification enables us to generate a set of simulated traces. Each trace is simply a sequence of API method invocation information. We set up a client to take in the simulated traces, and issue the corresponding method calls to interact with the component. The interactions are then recorded as program execution traces. We next invoke our miner to operate on this set of traces to produce an output specification. Lastly, we compare the output specification against the initial specification.

Our simulation model provides a controlled environment for experiments (*cf.*, [17]). Given a piece of software component under examination, we can alter the input API-interaction specification and later compare it with the corresponding output specification. This preliminary section provides a summary of our work on a quality assurance framework for *automaton-based* specification miners [18].

2.1 Specification Model

Input to our simulation is a specification model in the form of PFSA, a sample of which is shown in Figure 1. Based on this PFSA, we generate *simulated traces*. These simulated traces are then used to instruct the client to interact with the subject software, and the interactions are then recorded as execution traces. We refer the reader to Section 4 for a detailed description of this implementation.

Each node in the automaton represents a program state. There are four types of nodes: start, end, normal and error nodes. Each transition in the automaton represents a viable API method call from that state. For every transition, a probability will be attached to it. The probability attached to a transition indicates how likely the associated method call will be invoked from that source state. It is an invariant of any PFSA under consideration that all transitions emitting from a source, excluding error transitions (see the following paragraph), must have their probabilities summed up to 1.0.

The specification model can be “*injected with error*” by including error nodes and error transitions. Error transitions (drawn as dashed lines in Figure 1) model invalid API method calls from the state represented by its source node. This inclusion of error nodes and error transitions enable generation of erroneous traces; it aids the evaluation of the miner’s ability to learn in the presence of errors (*i.e.*, robustness). The allocations of error nodes and transitions will characterize the kind of errors allowed. Lastly, we do not assign any probability to error transitions, as we do not intend to micro-manage the generation of error traces.

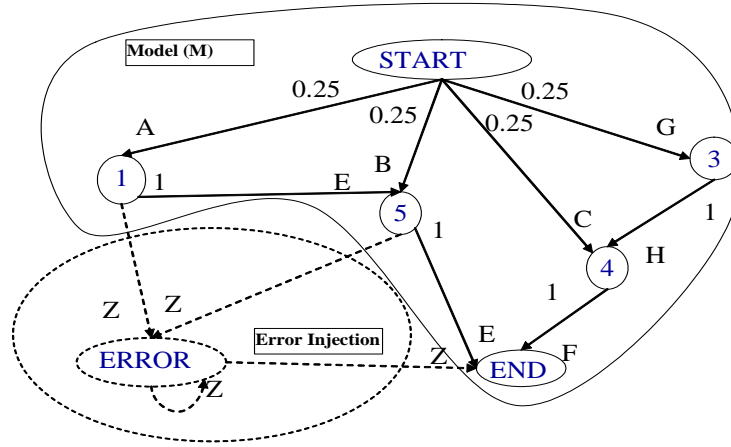


Fig. 1. Sample Simulator Model

Furthermore, large models (in terms of number of nodes or transitions) can be inputted to test the scalability of a miner. Hence, with injection of error and a variety of model sizes different levels of quality assurance can be obtained.

2.2 Simulated-Trace Generation

In order to explore the effect of mining under different initial automata, we simulate the interaction with the software from these automata.

An actual program trace can be mapped to a string over an alphabet as shown by Ammons *et al.* through a 'standardization' process [6]. Strings of symbols generated by our specification model (*aka.*, simulated traces) can be considered as an abstraction of actual program execution traces; *i.e.* a symbol representing a particular method call. Based on this abstraction, we generate simulated traces as strings of symbols.

Simulated-trace generation will generate two types of output - error, and normal traces. A *trace* is defined as a sequence of transition names that forms a path sourcing from the start node and sinking at the end node of a PFSA. We define an *error trace* as one that includes a transition sinking at an error node.

For each trace generated from a PFSA, we can determine its probability of being generated by multiplying together the probability of its constituents. We write $p(t)$ to denote the probability of a trace t .

To generate traces, we perform a *stratified random walk* guided by the probability of PFSA's transitions [19]. This ensures that highly probable traces (sentences) accepted by the PFSA model will statistically be more likely to appear in the multiset of generated traces. (We use the term "sentence" and "trace" interchangeably.)

Traces will continue to be generated until all transitions have been covered at least N times or *Max* number of traces have been generated. By adjusting the value N , we

can accommodate a slower learner that requires more than one trace from a specification to infer the automata model. By default we set N to 10 and Max to 10000.

Our algorithm is akin to the “code and branch coverage” criterion used in generating program test cases [20, 21]. Given a PFSA \mathcal{M} and a global percentage of error, our algorithm generates a multiset of traces T possessing the following property:

Property 1 *For a sufficiently large T , there is a $N > 0$ such that all (non-error) transitions in the PFSA \mathcal{M} occurs at least N times in the traces of T .*

This property ensures that all (non-error) transitions in M have the opportunity to be used for trace generation. The algorithm detail is depicted in Figure 2.

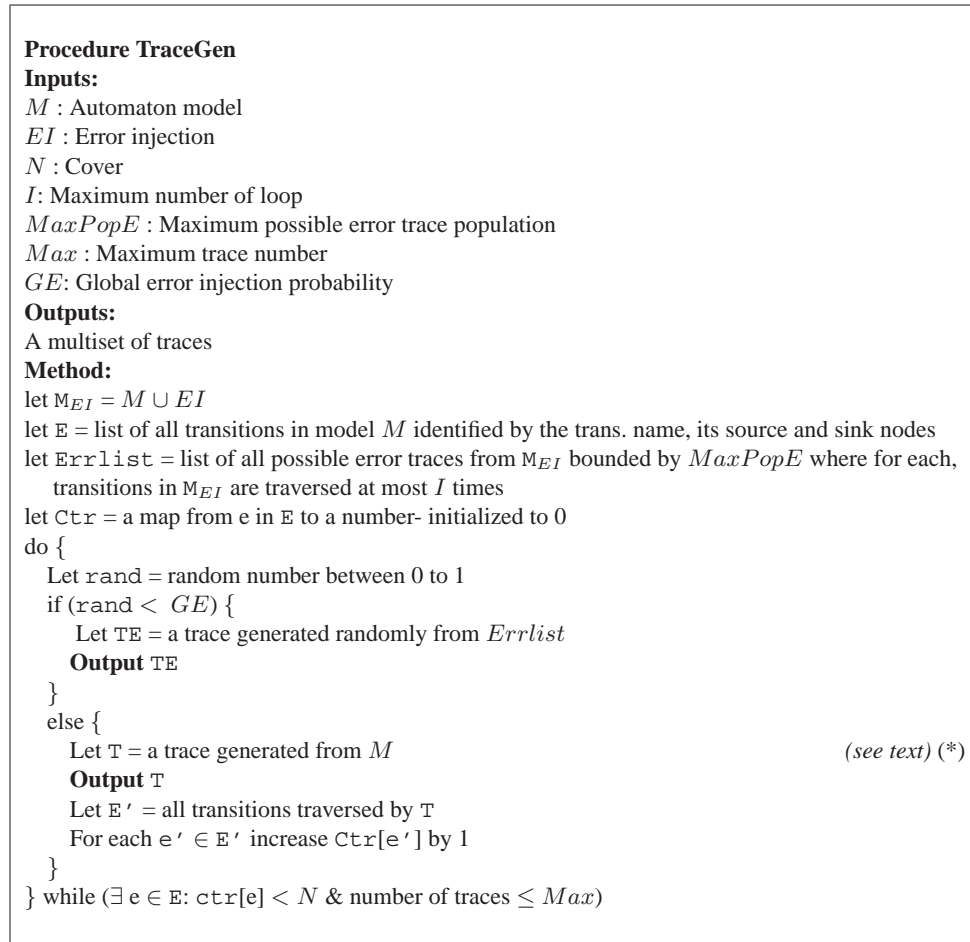


Fig. 2. Trace Generation Algorithm

M_{EI} is the PFSA M with error EI injected. At program point (*), A trace is generated by starting from start node of the model and independently ”throwing a dice” at each node for decision on which transition to take *according to the probability of the transitions* until end node is reached. Traces generated will then reflect the probabilities of the transitions in the simulator model (*i.e.* distribution of generated traces is governed by the model).

2.3 Precision, Recall and Co-emission

The terms *precision* and *recall* originated from the field of information retrieval, where they are defined as “*the proportion of retrieved documents which are relevant*” and “*the proportion of relevant documents retrieved*” respectively [22].

Analogously, traces can be considered as documents and automata as a pool/population of documents. Let the original automata be denoted by X and the inferred automata by Y . Precision and recall can then be defined as *the proportion of traces in Y that is accepted by X* and *the proportion of traces in X that is accepted by Y* where X is the original specification and Y is the inferred specification.

The total number of traces accepted by an automata can possibly be infinite. Hence precision and recall can only be statistically approximated; here, Property 1 ensures that the set of traces generated are statistically sound.

In the context of PFSA, a trace might possibly be generated by both X and Y , but their probability (of how frequently the trace will be generated) might differ greatly. Co-emission is therefore used to address this probabilistic concern.

Co-emission has been used in measuring similarity between two Hidden Markov Models. Lyngsø *et al.* propose several versions of similarity measurements [16]. One such metric which is adopted here, denoted by PS, provides an unbiased and normalized similarity measurement of two models X and Y :

$$PS(X, Y) = \frac{2 * P_{CE}(X, Y)}{(P_{CE}(X, X) + P_{CE}(Y, Y))}$$

$$P_{CE}(X, Y) = \sum_{s \in L(X \cap Y)} (P_X(s) P_Y(s)).$$

Here, $P_{CE}(X, Y)$ denotes a *co-emission probability*, determining the probability that a sentence s is generated by both X and Y independently. It measures how similar are the probabilities assigned to traces accepted by both X and Y . $P_X(s)$ and $P_Y(s)$ denote the probability of generating sentence s by X and by Y respectively.

3 SMArTIC Structure

SMArTIC aims to increase a miner’s precision, robustness and scalability by employing several novel techniques in specification mining. It leverages on the lessons learnt and experience accumulated from the past work done in this and related areas (*eg.*, [6], [4], [23], *etc.*). The success of SMArTIC hinges on the affirmation of the following two hypotheses:

Hypothesis 1 *Mined specifications will be more accurate when erroneous behavior is removed before learning than when they are removed after learning.*

Hypothesis 2 *Mined specifications will be more accurate when they are obtained by merging the specifications learned from clusters of related traces than when they are obtained from learning the entire traces.*

Hypothesis 1 is made from observing the system built by Ammons *et al.* [6]. In their work, a *coring* method is employed to remove erroneous transitions from the mined automaton. As this is performed on the output automaton, erroneous transitions are included during mining. Consequently, the performance of learning may be degraded. Moreover, pruning of transitions in an automaton may cause damage to the automaton, such as breaking an automaton into parts. This may then require substantial repairing of the automaton, and negate the effect of learning.

We believe that pruning of erroneous transitions should be done *before* learning. Consequently, we include a *filtering* process before the learning process in SMArTIC, as we shall describe in Section 3.1.

Hypothesis 2 is derived from the observation that the existence of unrelated traces may negate the effect of learning via generalization; *ie.*, they can lead to over-generalization. Therefore, by clustering related traces and performing learning on each cluster, the effect of inaccuracies in learning can be localized to within a cluster. We believe this will result in a more accurate mined specification. Consequently, we include a clustering process in SMArTIC, as we shall describe in Section 3.2.

The overall structure of SMArTIC is as shown in Figure 3. It comprises 4 major blocks, namely filtering, clustering, learning and merging blocks. Each block is in turn composed of several major elements. The filtering block filters erroneous traces to address the robustness issue. The clustering block divides traces into groups of 'similar' traces to address the scalability issue. The learning block generates specifications in the form of automata. The merging block merges the automata generated from each cluster into a unified one.

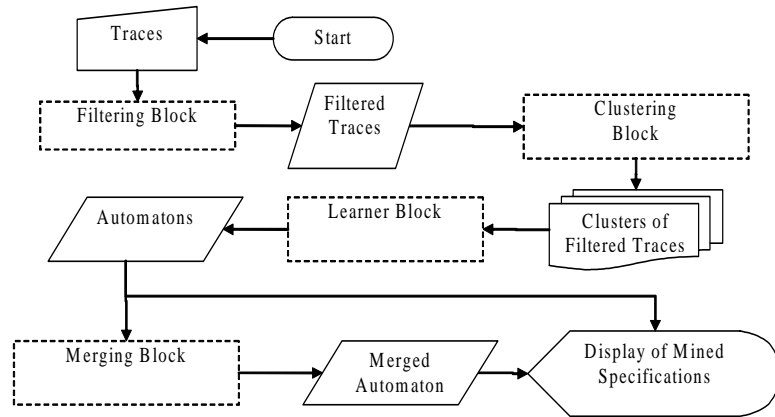


Fig. 3. SMArTIC Structure

3.1 Filtering Block

The filtering block aims to filter out erroneous traces based on common behavior found in a multi-set of program traces. To filter well, we need a representation of common behavior which is intuitive enough to be used for filtering. Since a trace is a temporal or sequential ordering of events, representing common behavior by “statistically significant” temporal rules will be appropriate. Certainly, temporal rules based on full set of temporal logics will be a good candidate, but it is desirable to have a more light-weight solution.

Given a set of traces, we would like to generate, through mining, rules of the form $\text{pre} \rightarrow \text{post}$, where both pre and post are sequences of alphabets occurring in traces. Semantically, such a rule has the following temporal interpretation: Given $\text{pre} = a_1 \dots a_m$ and $\text{post} = b_1 \dots b_n$, the temporal interpretation of $\text{pre} \rightarrow \text{post}$ is expressed in Linear Temporal Logic (LTL) notation [24] as

$$G(XF(a_1 \rightarrow \dots \rightarrow XF(a_m \rightarrow XF(b_1 \wedge \dots \wedge XF b_n))))$$

As an example, a rule $a \rightarrow bc$ asserts that at any trace point when a occurs, b must eventually occur after a , and c must also eventually occur after b .

There are two commonly used measures of “statistical significance” namely, *support* and *confidence* (c.f [25]). Support of a rule $\text{pre} \rightarrow \text{post}$ is the number of *trace points* exhibiting the property $\text{pre} \rightarrow \text{post}$. Confidence of the rule is the ratio of the number of *trace points* exhibiting the property $\text{pre} \rightarrow \text{post}$ to those exhibiting the property pre .

Rules having high confidence and reasonable support can be considered as “statistical” invariants. They thus characterize some general behaviors of a subgroup of traces. To detect outliers, only rules with high but less than 100% confidence will be useful. We call rules of $\text{pre} \rightarrow \text{post}$ format and exhibiting the above properties *outlier detection rules*.

Mined outlier detection rules will be used to filter out likely errors or unlikely behaviors. Any trace t_x of the following format $a_1 \dots a_i \dots a_{end}$ will be filtered out (as an outlier) by a rule-set RS iff the following holds:

$$\begin{aligned} \exists G(\text{pre} \rightarrow \text{post}) \in RS. \\ (\exists a_i, a_j. (1 \leq i \leq j) \wedge \neg(a_i \dots a_{j-1} \text{ satisfies pre}) \\ \wedge (a_i \dots a_j \text{ satisfies pre}) \\ \wedge \neg(a_{j+1} \dots a_{end} \text{ satisfies post})) \end{aligned}$$

The algorithm to filter out those traces that deviate from the general behaviors is depicted in Figure 3.1.

Implementation-wise, the structure of the filtering block is as shown in Figure 5. Outlier detection rules can be extracted efficiently by adding pre and post processing steps to a closed sequential pattern miner, BIDE [26]. The end result of the filtering block is a multi-set of filtered traces.

```

Procedure Filter
Inputs: Traces: A set of traces
           Rules: Outlier detection rules
Outputs:
           Filtered: A set of filtered traces
           Err: A set of traces deviant from rules
           FTRulesMap: A mapping from a trace to
                        a set of rules it satisfies
Method:
           Filtered, Err, FTRulesMap =  $\emptyset$ 
           For each trace  $t$  with format  $a_1 \dots a_i \dots a_{end}$  in Traces
             Let Satisfy =  $\emptyset$ 
             For each rule  $\langle pre, post \rangle$  in Rules
               For each  $a_i \dots a_j$  substring of  $t$  satisfying pre
                 If  $a_{j+1} \dots a_{end}$  satisfies post
                   Set Satisfy = Satisfy  $\cup$   $\{t\}$ 
                 else
                   Err = Err  $\cup$   $\{t\}$ 
                 break
             FTRulesMap.Add( $t$ , Satisfy)
             Filtered = Filtered  $\cup$   $\{t\}$ 
Output Filtered, Err, FTRulesMap

```

Fig. 4. Filtering Traces from Temporal Rules

Implementation Details. Sequential pattern mining takes as input SA (a set of sequences) and min_sup (minimum support level). It will then reports subsequences (or pattern) contained by at least min_sup no of sequences in SA. The number of such supporting sequences in SA is called ‘support’. Subsequences having support more than min_sup is called ‘frequent’. Sequential pattern miner will return both ‘frequent’ subsequences and its ‘support’.

Collected program traces, each of which is a sequence of method calls, can be considered as set of sequences SA. Inputting collected program traces to a sequential pattern miner will generate a set of subsequences (or patterns) of method calls that is supported by many traces.

A program often contain loops. Subsequence (or pattern) of method calls can be repeated more than once in a trace. Rather than counting the no of traces, we should count the no of locations within traces (*i.e. trace points*) where a pattern of method calls appear. ‘Frequent’ pattern of method calls should be based on on no of supporting *trace points* rather than no of supporting traces.

To facilitate counting of supporting *trace points* rather than traces, we perform pre-processing to input traces. Our pre-processing stage converts original multiset of traces (T_{Orig}) to its superset (T_{Result}) to facilitate approximate counting of *trace points* without changing the behavior of underlying sequential pattern miner. For every $t (t_0 \dots t_{end})$ in T_{Orig} ,

$$\{t_i \dots t_j | (t_i \dots t_j \text{ suffix_of } t) \wedge (\exists k. (k < i) \wedge t_k = t_j)\} \cup \{t\}$$

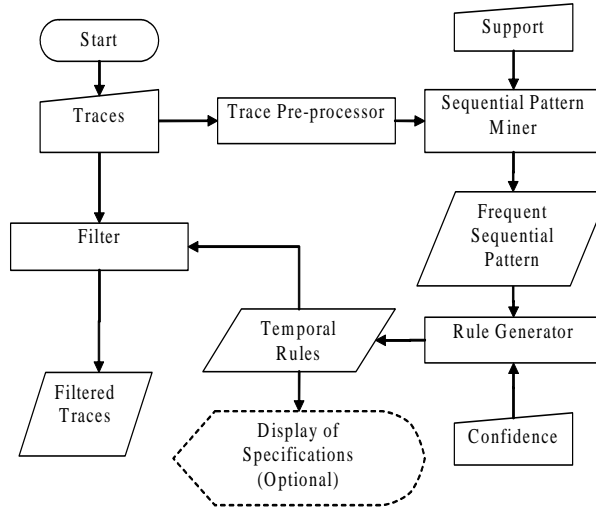


Fig. 5. Filtering Block Structure

is added to multiset T_{Result} . Inputting T_{Result} to a frequent sequential pattern miner will generate frequent pattern of method calls supported by a substantial number of trace points.

Any subsequence of a frequent subsequence will also be frequent (*i.e.* apriori property – *c.f.* [25]). Hence, a long frequent subsequence will generate a combinatorial number of frequent subsequences (*c.f.* [27]). Long frequent subsequences of method call might appear due to ‘wrapper’ effect of deep class hierarchy, be an effect of decomposition of a complex methods to a series of simpler ones, or correspond to initialization and termination of a protocol (*e.g.* connect \rightarrow login \rightarrow logout \rightarrow disconnect in FTP protocol).

To reduce the combinatorial number of frequent subsequences, closed sequential pattern miner have been proposed [27]. A closed sequential pattern, is a frequent pattern which is not a subsequence of another frequent pattern with the same support. A set of closed sequential pattern captures the same information as full set of sequential pattern without being combinatorial in number. BIDE [26] is an optimized miner of closed sequential pattern. Experimental results have shown its run-time is linear w.r.t. input size (no of sequences \times avg length of sequences). We run BIDE with preprocessed traces and support level Sup to generate a set of closed sequential patterns.

Frequent sequential patterns return by BIDE can be post-processed into pre \rightarrow post rules. Given two frequent sequential pattern $\langle A \rangle$ and $\langle A, B, C \rangle$ with support s_1 and s_2 , a rule $A \rightarrow BC$ can be generated with ‘confidence’ s_2/s_1 (*c.f.* [28, 25]). Given a subsequence f which is frequent, support (f) = $Max_{c \in Closed \wedge f \text{ subsequence-of } c} \text{support}(c)$. Hence, support of a frequent subsequence can be inferred from *Closed*. The conversion algorithm to generate rules from closed sequential pattern is as shown in Figure 6.

The algorithm receives as input a set of closed sequential patterns (*Closed*) and a minimum confidence level (*Conf*). It first build a prefix tree (or trie) of the closed sequential patterns. A prefix tree is a tree where there is one node for every common prefix. A closed sequential pattern can be considered a sequence of events. For each node $q \in \text{trie}$, $q.\text{event}$, and $q.\text{prefix}$ denotes its corresponding event and sequence of events from trie's root to q . A node's owner, $q.\text{owner}$, is the closed pattern sharing prefix $q.\text{prefix}$ that has the maximum count. The support of $q.\text{prefix}$ (denoted as $q.\text{count}$) is simply the count of $q.\text{owner}$. A closed pattern c 'shared' a trie node q , if $q.\text{prefix}$ is a prefix of c .

The trie is later traversed to locate 'interesting' nodes. A node q is interesting if, $\exists qd. ((qd \text{ child-of } q) \wedge (Conf \leq (qd.\text{count}/q.\text{count}) < 1))$. From such nodes, rules with confidence $\geq Conf$ but < 1 of the form $q.\text{prefix} \rightarrow \text{post}$, where post starts with $qd.\text{event}$ can be generated.

The algorithm will extract $\text{pre} \rightarrow \text{post}$ rules having confidence at least *Conf* but less than 100%. All rules generated will also be frequent since it is generated from closed sequential patterns. Some rules that can be inferred from others were pruned since they are redundant. The details of pruning and completeness issue are outside the scope of this technical report.

By trace pre-processing, rules mined will apply at every *trace points*. Hence, each rule can be expressed as $\text{pre} \rightarrow \text{post}$, where $\text{pre} \rightarrow \text{post}$ is of the following format $\text{pre}_1 \wedge XF(\text{pre}_2 \dots \wedge XF(\text{pre}_{end} \rightarrow (\text{post}_1 \wedge XF(\text{post}_2 \dots)))$ in LTL notation. These rules will have support of at least *Sup* and confidence $\geq Conf$ while $< 100\%$ – i.e. they are outlier detection rules¹. Hence, we have mined *outlier detection rules* efficiently by adding pre and post processing steps to a closed sequential pattern miner.

3.2 Clustering Block

Input traces might be "mixed up" from several unrelated scenarios, e.g. a group of related traces that represent a usage pattern of an API/component. Grouping unrelated traces together for a learner to learn might multiply the effect of inaccuracies in learning a scenario. Such inaccuracies can be further permeated into other scenarios through generalization.

The clustering block converts a set of traces into groups of related traces. Clustering is meant to *localize inaccuracies* in learning one sub-specification and prevent the inaccuracies from being permeated to other sub-specifications. Furthermore, by grouping related traces together, better generalization (*aka.*, less over-generalization) can be achieved when learning from each cluster.

Two major issues pertaining to clustering are: the choice of clustering algorithm and an appropriate similarity metric; i.e., measurement of similarity between two traces. The performance of the clustering algorithm is affected by appropriate similarity/distance metrics. Different clustering algorithms learn differently in terms of accuracy, efficiency

¹ Note, the confidence and support values of rules generated by our method are approximate figures.

```

Procedure Generate Rules
Inputs:
Closed: A set of closed sequential patterns (seq. of events)
Conf: Minimum confidence
Outputs:
Rules: Outlier detection rules
Method:
Step 1: Trie building and traversal
Let Trie = Build a prefix tree from Closed
           with each node representing an event
For each node p in Trie do
  Set p.prefix = Seq of events from root to p (inclusive) in Trie
  Set p.owner = f, where  $f \in \text{Closed} \wedge f \text{ shared } p$ 
                 $\forall g \in \text{Closed} \wedge g \text{ shared } p \cdot (f.\text{count} \geq g.\text{count})$ 
  Set p.postfix = Seq of events from p to leaf,
                 where seq from root  $\rightarrow p \rightarrow$  leaf
                 generates p.owner
  Set p.count = p.owner.count
end for
Let Interesting = {q | q  $\in$  Trie  $\wedge$ 
                   $\exists qd. ((qd \text{ child-of } q) \wedge$ 
                     $(Conf \leq q.\text{count}/qd.\text{count} < 1))$ }
Step 2: Rule generation
For each qi  $\in$  Interesting do
  Let qdesc = {qd | (qd child-of q)  $\wedge$ 
                 $(Conf \leq q.\text{count}/qd.\text{count} < 1)$ }
  Let RS = Generate rules pre  $\rightarrow$  post where,
            pre = qi.prefix  $\wedge$ 
            post qdesc.postfix
  Add RS to Rules
end for
Output Rules

```

Fig. 6. Rule Generation

and the level of user interaction required. (c.f. [25]) The general structure of the clustering block is as shown in Figure 7.

Clustering Algorithm We use a classical off-the-shelf clustering algorithm for our purpose, namely the k-medoid algorithm [29].² The k-medoid algorithm works by computing the distance between pairs of data items based on a similarity metric; this corresponds to computing the distance between pairs of traces. It then groups the traces with

² Another algorithm is K-means algorithm. It is not used here since we would be required to define the average/mean of a group of strings which might not be meaningful. Also, k-medoid algorithm has been found to be more accurate than k-means since it is less susceptible to outliers [25].

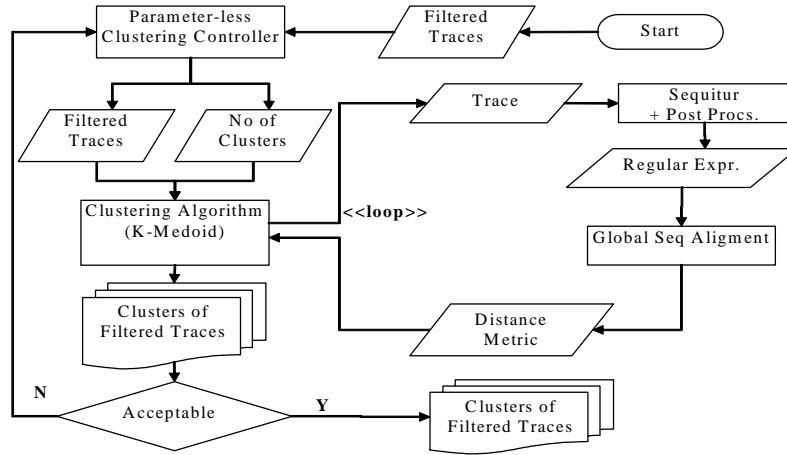


Fig. 7. Clustering Block Structure

small distances apart into the same cluster. The k in k-medoid is the number of clusters to be created.

In our implementation, we adapt the Turn* algorithm presented by Foss *et al.* [30] into the k-medoid algorithm. The Turn* algorithm can automatically determine the number of clusters to be created by considering the similarities within each cluster and differences among clusters. Our algorithm will repetitively increase the number of clusters. For each repetition, it will divide datasets into clusters and evaluate a measure of similarities within each cluster and differences among different clusters. The algorithm will terminate once a local maximum is reached. The algorithm is shown in Figure 8. At program point (*) a call is made to procedure CalculateSim defined in subsection 3.2.

Similarity Metric In many applications, comparisons between two data items are relatively clear – sometimes it only involves a simple subtraction of two numbers – *eg.*, (Avg. profit for company x) - (Avg. profit for company y), *etc.* However, a comparison of two program traces is neither so clear cut nor easily obtained.

Our first idea is to use *global sequence alignment* [31] to measure the distance between two traces. This alignment is frequently used to obtain similarity metrics of two DNA sequences. Its main idea is to insert “dash” or spaces within strings to obtain the most accurate matching of two strings. Different from the Knuth-Morris-Prat (KMP) algorithm [32], the sequence alignment algorithm finds the best approximated alignment(s) rather than the occurrence of an exact match. Alongside the best alignment(s), an overall similarity score will also be reported. We use this score as the similarity metric between the two program traces.

This first idea does not work well in practice because, contrary to normal strings, program traces exhibit some characteristics which make it difficult to measure similarity by a simple alignment of two traces. Specifically, a trace might only be different to

Procedure PLess–KMedoid**Inputs:**

IPTraces : Set of unique input traces to be split to clusters

Outputs:

Clusters: Set of clusters taking into account similarity within each cluster and differences among different clusters

Method:

Let *FeatureScoreList* = {}

Let *IsLocalMaxima* = false

Let *Clusters* = {}

For (k=1;k<#*IPTraces* && !*IsLocalMaxima*;k++)

 Set *Clusters* = Call *KMedoid* (*IPTraces*, k)

 Let *FeatureScore* = Call *CalculateScore* (*Clusters*)

FeatureScoreList.Add (*FeatureScore*)

 Set *IsLocalMaxima* = Check whether a local maxima has been found in *FeatureScoreList* by noting gradient change

end for

Output *Clusters*

Procedure CalculateScore**Inputs:**

IPClusters : Clusters of data items

Outputs:

Score: Feature score considering similarities within each cluster and differences among clusters

Method:

Let *MedoidCenters* = {*c*.*Medoid* | *c* ∈ *IPClusters*}

Let *MedoidCentersCluster* = Call *KMedoid* (*MedoidCenters*, 1)

Let *RepMedoid* = *MedoidCentersCluster*.*Medoid*

Let *SimMedoids* = 0.0

Let *SimWithinClusters* = 0.0

Step 1: Calculate similarities between clusters medoids

For each *medoid* in *MedoidCenters* not = *RepMedoid* do

 Set *SimMedoids* + = *CalculateSim* (*MedoidCenters*[*i*], *RepMedoid*) (*see text*)(*)

end for

Average *SimMedoids* by #*MedoidCenters*

Step 2: Calculate similarities within each clusters

For each *cluster* in *IPClusters* do

 Let *AClusterSim* = 0.0

 For each *dataitem* in *cluster* != *cluster*.*Medoid*

 Set *AClusterSim* + = *CalculateSim* (*dataitem*, *cluster*.*Medoid*) (*see text*)(*)

 end for

 Average *AClusterSim* by #*cluster*

 Set *SimWithinClusters* + = *AClusterSim*

end for

Average *SimWithinClusters* by #*IPClusters*

Let *Score* = *SimWithinClusters* – *SimMedoids*

Output *Score*

Fig. 8. PLess KMedoid Algorithm

another due to different numbers of loop iterations during program execution. As an example, consider the following program segment:

```
function APICLIENT_ABCD (outer_iter,
inner_iter[]) {
  for (int j=0;j<outer_iter;j++) {
    int k=0; Call API.A ();
    do{ k++;
      Call API.B();
      Call API.C();
    }while (k<inner_iter[j]);
    Call API.D ();
  }
}
```

Suppose that APICLIENT_ABCD is a client function of an API. It is conceivable that the API interaction patterns for various runs via the function call APICLIENT_ABCD with different input parameters should be grouped together. So, for a run with parameters `outer_iter = 2` and `inner_iter = [2, 3]`, the generated trace is ABCBCDABCBCBCD. For another run with `outer_iter = 1` and `inner_iter = [1]`, the generated trace is ABCD. Now, if we simply align these two strings, even in their best alignment their similarity score will be too low for them to be grouped into the same cluster.

Our solution to the above problem is to instead compare the regular expression representations of the two traces rather than their actual sequence of alphabets. Converting to its regular expression, the first trace will be $(A(BC)+D)^+$ which corresponds closely to ABCD. Note, we choose to only use “+” quantifier (*ie.* one or more repetitions) but not “*” quantifier since every expression can be reduced to $((a_1)^* (a_2)^* \dots (a_n)^*)^*$, where $\forall_{1 \leq i \leq n} a_i \in \text{Alphabet}$.

We obtain the regular expression representation by converting a trace to its hierarchical grammar representation using Sequitur [33]. The output of sequitur will be post-processed to construct the regular expression representation and then be fed in as input to global sequence alignment. With these we obtain a method to find a reasonable distance metric for the similarity of program traces. Generated sequitur grammar of a trace will be processed through several passes of reduction to produce a regular expression representation. The distance calculation algorithm and grammar reduction to regular expression algorithm are shown in Figure 9. It makes use of HieUni algorithm that is shown in Figure 10.

3.3 Learning Block

Although temporal rules have also been used ([34]) to capture certain information of a program specification, automata have been commonly used in capturing specifications, especially protocol specifications. The purpose of this learning block is to learn automata from clusters of filtered traces. This block is actually a placeholder in our architecture. Different PFSA specification miners can be placed into this block, as long as they meet the input-output specification of a learner. Once a learner is plugged in, it will be used to mine the traces obtained from each cluster. At the end, the learner produces one mined automaton for each cluster.

In the current experiment, we choose to use a PFSA specification miner that has been used for software specification mining earlier, *ie.* sk-strings learner [35].

```

Procedure CalculateSim
Inputs:
Trace1 : Program trace 1
Trace2 : Program trace 2
Outputs:
SimScore: Similarity between program traces
Method:
Let Regex1 = Call StringToRegEx (Trace1)
Let Regex2 = Call StringToRegEx (Trace2)
Let SimScore = Global sequence alignment score between Regex1 and Regex2.
Output SimScore

```

```

Procedure StringToRegEx
Inputs:
Input : Input string
Outputs:
RegEx: Regular expression
Method:
Let rOrig = Rules from invoking Sequitur(Input)– format: var->(terminal|var)+
Let rMap = Mapping btw left (r[i].L) to right part (r[i].R) of rules
Let oR = Ordered rules where oR[i].R.Contains(oR[j].L)  $\iff$  (j <= i)
Let RegEx = oR[0].R
Let varlst = Get vars in RegEx.R ordered by oR
while (varlst != empty) do
  Let r0 = varlst[0]
  Set RegEx = Call HieUni (RegEx)
  Set RegEx = Replace r0.L with r0.R in RegEx
  Set varlst = Get vars in RegEx.R ordered by oR
Set RegEx = Call HieUni (RegEx)
Let block = set of first level ()+ constructs in RegEx
Let fvMap = {}
for each (b)+ in block starting from longest b do
  Let bplusset = indices to all (b)+ in RegEx
  Let bset = indices to all b in RegEx but not in bplusset
  Let fvar = a free variable
  Set RegEx = replace all b indexed in bset in RegEx with (b)+
  Set RegEx = replace all (b)+ in RegEx with fvar
  fvMap.Add (fvar,(b)+)
if (Input != RegEx)
  Call StringToRegEx (RegEx)
for each (fvar , content) in fvMap do
  Replace fvar in RegEx with content
Output RegEx

```

Fig. 9. Trace Similarity Score Algorithm

Sk-strings learner is used by Ammons *et al.* to mine the specification of the X11 windowing library [6]. It is an extension of the k -tails heuristic algorithm of Biermann and Feldman [36] for learning stochastic automata. In k -tails, two nodes in a

```

Procedure HieUni
Inputs:
  Input : Input string
Outputs:
  Unified: Input string with substrings unified
Method:
  Let pT = Convert Input to tree of seq (i.e. ()+ construct) and term
  Let sot = Get seqs in pT containing only terms
  while sot != empty do
    for each seq in sot do
      Let term[0] = first child of seq
      Let reprMerged = term[0].Text
      Let bSeq = F
      for each term[i] child of seq; i >= 1
        if (term[i] != term[i-1])
          if (bSeq)
            Set reprMerged += "("+term[i-1].Text+"+"
          end if
          Set reprMerged += term[i].Text
        else Set bSeq = T
        end if
      end for
      Let newT = Create new term with text reprMerged
      if (seq != pT.Root)
        Replace seq with newT from seq.Parent
      else Replace seq with newT as pT.Root
      end if
    end for
    Set sot = Get seqs in pT containing only terms
  end while
  Let Unified = pT.Root.Text
Output Unified

```

Fig. 10. Hierarchical Unification

constructed automata are checked for equivalence by looking at subsequent k -length strings that can be generated from them. Different from k -tails, in sk -strings, subsequent strings need not necessarily end at an end node, except for strings of length less than k . Furthermore, only the top $s\%$ of the most probable strings that can be generated from both nodes are considered. Implementation-wise, the sk -strings learner first builds a canonical “machine” similar to a prefix-tree acceptor from the traces. The nodes in this canonical machine are later merged if they are indistinguishable with respect to the top $s\%$ of the most probable strings of length at most k that can be generated starting from them.

3.4 Merging Block

The merging process aims to merge multiple PFSA's produced by the learner into one such that there is no loss in precision, recall and likelihood before and after the merge. Equivalently, the merged PFSA accepts exactly the same set of sentences as the combined set (*i.e.*, *union*) of sentences accepted by the multiple PFSA's.

The primary purpose of the merging process is to reduce the number of states residing in the output PFSA by collapsing those transitions behaving "equivalently" in two or more input PFSA's, thus improving scalability.³

Merging of two PFSA's involves identification of *equivalent* transitions between the two PFSA's. The output PFSA contains all transitions available in the input PFSA's, with each set of equivalent transitions represented by a single transition. Two identically-labelled transitions from two different PFSA's are considered *equivalent* if one the following conditions holds: (1) Both their source nodes share the same set of suffixes, (2) Both their sink nodes share the same set of prefixes.

Given a node n in a PFSA and a string accepted by PFSA that involves a transition, δ say, emitted from n , we define a suffix of n with respect to the string as the suffix of that string beginning with δ . Similarly, we define a prefix of n as the prefix of that string ended just before δ . For instance, given the transition $n_1 \xrightarrow{\delta} n_2$, and a string

$$t_1 t_2 \cdots t_{m-1} \delta t_{m+1} \cdots t_p,$$

A suffix of n_1 is the string $\delta t_{m+1} \cdots t_p$, and a prefix of n_1 is the string $t_1 t_2 \cdots t_{m-1}$.

In the case where no transition emitting from n appears in the string, both the prefix and suffix of n with respect to that string are just null.

Extending from the definitions above, the set of prefixes/suffixes of n in a PFSA is the set of prefixes/suffixes of n with respect to all strings accepted by the PFSA.

The definition of equivalent transitions above admits closure property: If two transitions are equivalent as defined by sharing the same set of suffixes, then each of the transitions in the suffix-set is also an equivalent transition. The same behavior can be observed from equivalent transitions sharing the same set of prefixes.

The Merging process also ensures that the likelihood of traces generated by the output PFSA remains the same as that of the combined input PFSA's. More specifically, let \mathcal{A} be the output PFSA and \mathcal{A}_i ($i = 1..n$) be the input PFSA's. Let δ be a transition in \mathcal{A} , then

$$p_{\mathcal{A}}(\delta) = \sum_{i=1}^n w_i \cdot p_{\mathcal{A}_i}(\delta)$$

where $p_{\mathcal{A}}(\ell)$ and $p_{\mathcal{A}_i}(\ell)$ represent the probabilities of the transition ℓ located in PFSA \mathcal{A} and its equivalent transitions occurring in \mathcal{A}_i . $p_{\mathcal{A}_i}(\ell)$ is assigned to 0 if no equivalent transitions occur in \mathcal{A}_i . w_i is the weightage given to each cluster hosting their own input PFSA; it is the ratio of the number of traces in the cluster to the number of total traces in the entire system.

Implementation-wise, the closure property of equivalent transitions enables an incremental detection of such transitions, starting from either the start node (for finding

³ A secondary objective is to enable comparison of structural similarity between the mined PFSA and the original PFSA. We omit the detail in this paper.

prefixes) or the end node (for finding suffixes) of a PFSA.

Implementation Details Merger block will merge PFSA's produced by learner block to a unified one. The merging is performed in the "safe" way where no further generalization is performed during the merging process. Hence the set of language accepted by the final PFSA will be the union of the set of languages accepted by each of the PFSA's. The merging process is performed iteratively by merging 2 PFSA's at a time. Let's call them X and Y for ease of reference. The algorithm for automaton merger is as shown in Figure 11.

There are 6 steps in automaton merge algorithm. They are (1). Handling of exception cases, (2). Creation of unifiable list, (3). Creation of mergable list, (4). Creation of equivalence class, (5). Creation of structural merge and (6). Creation of full merge.

We consider it an exception case if there is a transition sinking in start node. This is such since the algorithm assume that the start nodes of the two automaton's are unifiable/mergable. However, if one automaton has start node as sink node of a transition and not the other, the two start nodes are not unifiable. In step 1, if this exception case happens a new start node is added with ϵ -transition to the original start node.

Two nodes of X and Y are considered unifiable if they shared common set of prefixes. In step 2, we create a list of pairing between $x \in X$.Nodes to $y \in Y$.Nodes where x and y are unifiable (*i.e.* creation of unifiable list). Two nodes of X and Y are considered mergable if they shared common set of suffixes. In step 3, we create this list of pairings (*i.e.* creation of mergable list).

The two lists are created by generating and solving constraints in a similar way with only minor differences. Unifiable list is created by unifying starting from start node top-down while mergable list is created by unifying starting from end node bottom-up. To check whether $xNode$ and $yNode$ are unifiable, only transitions having $xNode$ or $yNode$ as their source (*i.e.* $x.Next$ and $y.Next$) need to be considered. Vice versa, only transitions having $xNode$ or $yNode$ as sink (*i.e.* $x.Prev$ and $y.Prev$) need to be considered for $xNode$ and $yNode$ to be mergable.

A node can be unifiable to more than one node. In step 4, we would like to create these equivalence classes. A representative node is created for each equivalence class. Since two mergable nodes shared a common set of suffixes, it must be the case that all their descendants (*i.e.* nodes reachable from them) must be mergable as well. Hence we can create transitions between representative nodes. The end result is a sub-automaton at the bottom of the merged automaton.

In step 5, we create a partially merged automaton by traversing from root node according to unifiable list created in step 2. In this step, we create the sub-automaton at the top of the merged automaton.

In step 6, we augment this partially merged automaton by adding in nodes and transitions in X and Y that has not been included to the merged automaton. If a transition leads to a mergable node m , we add a transition to the m 's representative node in sub-automaton created in step 4 – no descendants of m need to be traversed anymore.

After the above 6 steps, the result is a merged automaton which is "safe" where no additional generalization was made during the merging process. The algorithm refers to unspecified procedures. These procedures detail can be found in the Appendix.

```

Procedure AutomatonMerge
Inputs:
X : First automaton to merge
Y : Second automaton to merge
TrainSetX : Set of traces to train X
TrainSetY : Set of traces to train Y
Outputs:
Merged: Merged automaton with  $L(Merged) = L(X) \cup L(Y)$ 
Method:
Let WX =  $\frac{|TrainSetX|}{(|TrainSetX|+|TrainSetY|)}$ 
Let WY =  $\frac{|TrainSetY|}{(|TrainSetX|+|TrainSetY|)}$ 
Let XUsed = {}
Let YUsed = {}
Let TransXUsed = {}
Let TransYUsed = {}
Let SharedNodes = {}
Let RepObjSet = {}
Step 1: Handle Exception Cases
If (#X.StartNode.Prev != 0) X.AddTempStartNode()
If (#Y.StartNode.Prev != 0) Y.AddTempStartNode()
Step 2: Create Unifiable List
Let Uni = Call CreateUnifiableList(X,Y,T)
Step 3: Create Mergable List
Let Merge = Call CreateUnifiableList(X,Y,F)
Step 4: Create Equivalence Class
Let DSSet = Call CreateEqClass (Merge,RepObjSet)
Step 5: Create Structural Merge
Let root = Create new node
Call CreateStructuralMerge
(X.Start,Y.Start,root,WX,WY,Uni,XUsed,YUsed,TransXUsed,TransYUsed)
Set SharedNodes = {n|(n,d) ∈ RepObjSet}
Step 6: Create Full Merge
Call CreateFullMerge (X.Start,DSSet,Empty List,SharedNodes, WX,XUsed,TransXUsed)
Call CreateFullMerge (Y.Start,DSSet,Empty List,SharedNodes, WY,YUsed,TransYUsed)
Let Merged = New automaton with root in root
Output Automaton Merged

```

Fig. 11. Automaton Merger

4 Jakarta Commons Net

Jakarta Commons Net [14] is a set of reusable open source java code implementing the client side of many commonly used network protocols. We built a simple CVS (Concurrent Versions System) functionality on top of an FTP library provided by Jakarta Commons Net.

4.1 Protocol for CVS-FTP API Interaction

This CVS functionality can be considered a client of Jakarta Commons Net with a certain protocol pattern. Our CVS class and Commons Net library can be instrumented to generate traces which were then inputted to SMArTIC and sk-strings. The resultant models are then compared with the original CVS specification to evaluate the feasibility of SMArTIC in improving the accuracy of the results over the sk-strings learner.

There are six common FTP interaction scenarios in our CVS implementation: Initialization, multiple-file upload, download, and deletion, multiple-directory creation and deletion. All scenarios begin by connecting and logging-in to the FTP server. They end by logging-off and disconnecting from the FTP server. The client side only maintains a record of files backed-up in the FTP server.

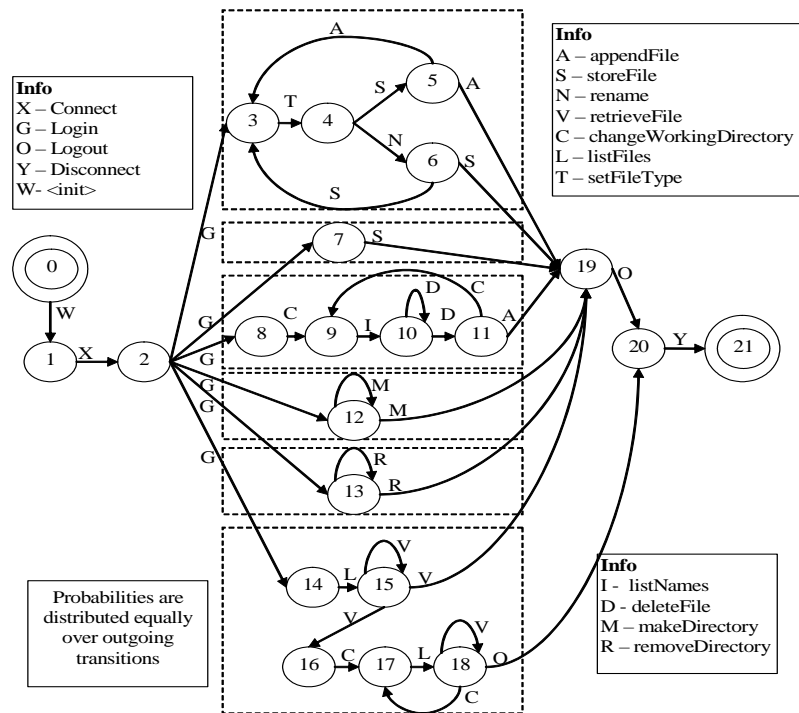


Fig. 12. CVS Protocol

All these scenarios are depicted in the automata shown in Figure 12. The dashed boxes, from top to bottom, represent upload files, initialization, delete files, make directories, remove directories and download files scenario respectively.

4.2 Instrumenting Jakarta Commons Net

We instrument Jakarta Commons Net with JRat, the Java runtime analysis toolkit [37]. The instrumentation is composed of an instrumentation byte code injection and a trace collection part. By default, JRat logs execution traces by associating them with a localized context. This context is simply a list of method calls in the runtime stack (*ie.*, `main() -> FTPClient.<init> -> TelnetClient.<init>`). Information having the same context is grouped together. Given a class file to instrument, JRat will add instrumentation code to all methods except the constructor.

We modified JRat core classes and added a plug-in to it. The following features were added:

Capturing order of method calls along with context. We would like to capture the temporal order of method calls together with the context. However, JRat may destroy the order of method calls in the context. Information on calls to a method at two different times under the same context but with different temporal ordering should not be grouped together.

Instrumentation of class constructor. In order to capture the hierarchy of method calls well, we need to instrument the class constructor as well. The class constructor might call other methods. We would like to capture the information that their context are the same but different from method calls called at constructor context.

Capturing object identities. Our plug-in captures object identities through their hash-code as well as the original class names where they are instantiated from (*ie.*, rather than the parent class where a called function is defined). The JRat structure provides good support for these information.

Thread slicing and Scalability. We slice traces into threads and generate separate trace file for each of them. For scalability, no large trace related information are stored in memory. They are always output incrementally.

The result of JRat instrumentation is an injection of tracing byte code into java classes. Running the instrumented code will produce a tree of method calls capturing their order and context represented as an XML document. Methods called earlier will print earlier in the XML document, each method will have its context as its ancestors in the XML tree.

4.3 Trace Collection and processing

We construct a wrapper class that takes in the automata shown in Figure 12 and invokes `org.apache.commons.net.ftp.FTPClient` accordingly. The wrapper class will generate a list of sequences of method invocations by traversing the automata from the start to the end node multiple times until coverage criteria is met. The result is a simulation of a regression testing of CVS-FTP API interaction.

Each invocation of a method of `FTPClient` may generate exceptions, especially `FTPConnectionClosedException` and `IOException`. Hence the code accessing the `FTPClient` methods need to be enclosed in a `try..catch..finally` block. Every time such an exception happens we simply logout and disconnect from the FTP server. This is simulated by adding error transitions as shown in Figure 13. Ten percent

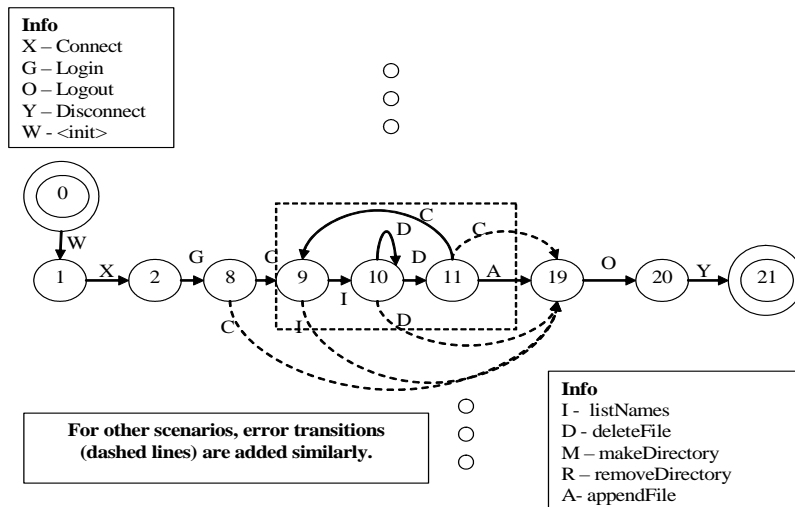


Fig. 13. CVS Protocol With Error

error is assumed and erroneous trace will be injected to 10% of the generated list of sequences of method invocations.

The trace file generated is likely to be huge because of the wrapper effect and long class hierarchies. On the other hand, what we really need are the traces capturing interaction between our own CVS classes and `FTPCClient`. To get that, we process the trace file as follows : (1) Traverse the XML trace file depth-first, and locate all the first invocations of the client method calls. Each such location will correspond to a scenario trace sequence. (2) From the above locations, traverse depth-first, and locate all the first invocations of the API method calls. The API method calls might not be directly below the client method calls in the trace file XML hierarchy.

4.4 Protocol Specification Generation and Results

The collected traces are inputted to different miners: SMArTIC (with sk-strings in the learner block), SMArTIC without filtering, SMArTIC without clustering, sk-strings with coring and standalone sk-strings. The coring threshold is set at 0.2 level. SMArTIC filtering confidence and support is set at 0.8 and 0.1 respectively. Default parameter settings are used for sk-strings both when standalone and within SMArTIC.

A protocol specification is then produced and compared against the original one (as shown in Figure 12) in terms of precision and recall. We repeat the above experiments 100 times using different lists of scenario trace sequences.

The following table shows the results of our experiment. The columns **Precs**, **Recall** and **PS** correspond to precision, recall and unbiased, normalized co-emission respectively (as defined in subsection 2.3).

	Precs	Recall	PS
SMArTIC	0.484	0.981	0.653
SMArTIC w/o filtering	0.426	1	0.616
SMArTIC w/o clustering	0.263	0.984	0.532
sk-strings(coring)	0.289	0.581	0.447
sk-strings	0.225	1.000	0.533

As shown in the table, SMArTIC improves the precision and co-emission while maintaining good recall of CVS protocol inference. The precision of SMArTIC are more than double the precision of sk-strings.

Both filtering and clustering help in increasing precision while maintaining good recall and equivalent or even better co-emission.

The drawback of coring is shown in the results where recall drops by almost half. Although precision is increased, there is a heavy penalty in recall: Pruning erroneous behaviors unavoidably removes a significant proportion of correct behaviors.

It is also of interest to know the number of erroneous traces our filtering algorithm filters out. On the average it filters out 43% of erroneous traces while only 4% of valid ones.

We have conducted thorough experiments using this application to verify both our hypotheses. Evaluating the statistical significance of our experiment results, we perform one-tailed t-test on the experiment results. The result is shown in table below.

H0	Result	Conf L
mr(SMArTIC w/o clus) = mr(cor)	R	99.99%
mp(SMArTIC w/o clus) = mp(cor)	NSE	90%
mps(SMArTIC w/o clus) = mps(cor)	R	99%
mp(SMArTIC w/o fil) = mp(str)	R	99.99%
mr(SMArTIC w/o fil) = mr(str)	NSE	90%
mps(SMArTIC w/o fil) = mps(str)	R	99%

mr,mp and mps denotes mean recall, mean precision and mean unbiased, normalized co-emission respectively. SMArTIC w/o clus, SMArTIC w/o fil, cor and str denotes SMArTIC without clustering (only filtering on), SMArTIC without filtering (only clustering on), sk-string with coring and sk-string (standalone) respectively. R and NSE denotes rejected and not sufficient evidence respectively.

The first result shows that it is statistically significant that mean recall of SMArTIC w/o clustering (only filtering on) is higher than mean recall of coring. The second result shows that mean precision of SMArTIC w/o clustering is not statistically different from mean precision of coring. The third result shows that co-emission of SMArTIC w/o clustering is higher than co-emission of coring. Hence, mined specification will be more accurate (improved recall and co-emission with equivalent precision) when erroneous behavior is removed before learning (in SMArTIC w/o clustering) than when they are removed after learning (in sk-string (with coring)). This supports our first hypothesis (stated in Section 3).

The fourth result shows that it is statistically significant that mean precision of SMArTIC w/o filtering (only clustering on) is higher than mean precision of sk-string

(standalone). The fifth result shows that mean precision of SMArTIC w/o filtering is not statistically different from mean precision of sk-string (standalone) – precision of both are exactly the same with mean 1 and standard deviation 0. The sixth result shows that co-emission of SMArTIC w/o filtering is higher than co-emission of sk-string (standalone). Hence, mined specification will be more accurate (improved precision and co-emission with equivalent recall) when it is obtained by merging the specifications learned from clusters of related traces than when it is obtained from learning the entire traces. This supports our second hypothesis (stated in Section 3).

5 Further Experiments

The experiment with CVS specification in Section 4 provides positive evidence that SMArTIC is a feasible architecture for improving mining accuracy; it also provides strong evidence to support our hypotheses stated in Section 3.

In this section, we perform further experiments on SMArTIC, not just on its accuracy, but also on its robustness and scalability. To this end, we have conducted almost 2000 experiments to support the superiority of SMArTIC.

Our experiments use the same set of miners, with sk-strings learner being employed either in standalone mode or in co-operation with other processes, especially as the learner block of SMArTIC.

5.1 Material

In the first set of experiments, two sets of sub-experiments using different types of error injection were performed to evaluate the two learners' performance in terms of robustness. These experiments are performed on sk-strings (standalone), sk-strings (with coring) and SMArTIC. For SMArTIC case, we disable the clustering sub-system to measure the effect of the filtering block.

We simulated the automaton generated by Ammons *et al.* in their analysis of the X11 windowing library (*cf.* [6]) – as shown in Figure 14. However, we modified the model slightly so that it was *without any non-determinism and repeated use of alphabet assigned to transitions* and we added probabilities. Probabilities are *distributed equally* to transitions from the same source node (not shown in the diagram). This is meant to produce a base model that can be learned perfectly. Error nodes and transitions were then injected to the automaton to conduct robustness tests. The models with different injections of errors (nodes and transitions labelled as Z and shown with dashed lines) are shown in Figure 15(a) & (b). Each of the two types of injections of errors shown in Figures 15(a) and (b) respectively corresponds to a separate set of sub-experiments.

We expect the specification miners to be able to filter out errors. We compared the inferred automaton with the simulator model shown in Figures 15(a) and (b) without error nodes and transitions and recorded the similarity and difference metrics. We generated traces using trace generation algorithm (briefly mentioned in subsection 2.1) and capped the maximum number of traces generated to 10,000. Four, eight and ten percent levels of error were injected to the system (*ie.*, 4, 8 and 10 percent of generated traces respectively will be erroneous). We assume the error level is unknown to the learner




Fig. 14. X11 Windowing Library Model




Fig. 15. Models of Specification with Error

except that it is low. Hence, the threshold used for coring was set to 0.2. SMArTIC's filter confidence is also set at an equivalent level of 0.8 while its support is set at 0.1. In each case, we ran a hundred experiments and recorded the average performance.

In the second experiment, we evaluated the scalability of the learners by generating distinct models of various sizes. Two sets of sub-experiments were conducted, each with a different independent variable. In the first set, we varied the number of nodes (by 10, 20, 30 and 40) in the specification model and maintained the number of outgoing transitions per node to at most four. In the second set, we varied the number of outgoing transitions per node (by 3, 5, 7 and 9) and maintained the number of nodes at 10. For each case, we performed 10 experiments and recorded their average performance.

We automatically generated distinct models having n nodes and a maximum of m transitions per state with common start and end nodes. Transition labels were chosen from a pool of a fixed number of labels randomly. Loops were introduced based on the principle of locality where loops between child and parent/ancestor nodes (including self-loop) occur with higher probability than those connecting to distant sibling nodes. The above properties are meant to generate *reasonably* complex models that are more likely to mimic reasonable protocols even in a large system (eg., business logic of an enterprise application).

Our model generation algorithm is shown in Figure 16. The algorithm initially create a connected automata structure in the form of a tree until N nodes have been created. Next, depending on loop level and locality level, a set of additional transitions will be introduced creating possibly loops and adding complexity to the model. Loop level and locality level by default are set to 0.4 and 0.8 respectively.

```

Procedure ModelGen
Inputs:
 $N$  : Total number of nodes,  $> 0$ 
 $M$  : Max number of transitions per node
 $Labels$  : Pool of labels
 $P_{Loop}$  : Loop level (from 0 to 1)
 $P_{Ancestor}$  : Locality level : Probability of loop to ancestor (including self)
Outputs:
 $Model$ : Model having  $N$  nodes with each node having at most  $M$  trans with labels from  $Labels$ .
Method:
Let  $Root$  = Create a new node for root node of model
Let  $NodeList$  = List of nodes so far, init to  $\{Root\}$ 
Let  $WorkList$  = List of nodes to process, init to  $\{Root\}$ 
// Step 1: Create Connected Automata Structure
while ( $\#NodeList < N$ ) do
  Let  $curNode$  = A random node from  $WorkList$ 
  Let  $transNo$  = A random no btw 1 to  $M$ 
  Set  $curNode.maxTrans$  to  $transNo$ 
  for ( $i = 0; i < transNo \ \&\& \ \#NodeList < N; i++$ )
    Let  $newNode$  = Create a new node
    Add transition from  $curNode$  to  $newNode$  with label randomly pick from  $Labels$ 
    Add  $newNode$  to  $NodeList$  and  $WorkList$ 
    if ( $i == \#transNo - 1$ )
       $WorkList.Remove(curNode)$ 
  end for
end while
// Step 2: Add Loop/Extra Transitions
Let  $NodeLoopList$  = Take  $P_{Loop} * \#WorkList$  nodes randomly frm  $WorkList$ .
for each node  $NLoop$  in  $NodeLoopList$ 
  Let  $transNo = NLoop.maxTrans$  undefined? A random no btw 1 to  $M: NLoop.maxTrans$ 
  while ( $NLoop$  has less than  $transNo$  trans.)
    Add a loop to  $NLoop$ 's ancestors with prob.  $P_{Ancestor}$  or to siblings otherwise
  end while
end for
Output Model with root node set to  $Root$ 

```

Fig. 16. Model Generation Algorithm

These experiments were performed on sk-strings (standalone) and SMArTIC. In the SMArTIC case, we measure the effect of the clustering block by disabling the filtering sub-system.

We generated traces and capped the maximum number of traces at 10,000. No error was injected to the system. Since we imposed a cap of 10,000 traces, there might be a concern that training traces might not satisfy the coverage criterion for a model of large size. This was not the case in our experiments, as only once was the cap reached; for the other 159 experiments, the coverage criterion was met first.

5.2 Experiment 1 Findings

Here, we evaluated the robustness of sk-strings, sk-strings (coring) and SMArTIC with two different injections of errors. The models with different injection of errors are shown in Figures 15(a) and (b).

Results. The experiment results for ErrModel1 and ErrModel2 are captured below. Column E% corresponds to the percentage of erroneous traces. Columns Precs, Recall and PS correspond to precision, recall and unbiased, normalized co-emission respectively (as defined in subsection 2.3).

Error Model 1							
sk-strings				sk-strings(coring)			
E%	Precs	Recall	PS	E%	Precs	Recall	PS
4	0.946	1.000	0.946	4	0.999	0.823	0.864
8	0.908	1.000	0.948	8	0.998	0.828	0.867
10	0.883	1.000	0.950	10	0.990	0.845	0.875
SMArTIC							
E%	Precs	Recall	PS	E%	Precs	Recall	PS
4	0.999	1.000	0.946	10	0.981	1.000	0.948
8	0.993	1.000	0.946				
Error Model 2							
sk-strings				sk-strings(coring)			
E%	Precs	Recall	PS	E%	Precs	Recall	PS
4	0.947	1.000	0.947	4	0.829	0.962	0.863
8	0.898	1.000	0.947	8	0.812	0.909	0.848
10	0.875	1.000	0.948	10	0.816	0.886	0.849
SMArTIC							
E%	Precs	Recall	PS	E%	Precs	Recall	PS
4	0.994	1.000	0.946	10	0.974	1.000	0.949
8	0.986	1.000	0.946				

Analysis. For sk-strings, all traces generated by the simulator model X were accepted by the inferred model Y . On the other hand, we noted a drop in the acceptance of traces generated by Y . This drop is slightly larger to the noise injected ((5.4%,5.3%) vs. 4%, (9.2%,10.2%) vs. 8% and (11.7%,12.5%) vs. 10%); learner precision degrades in the presence of erroneous traces. We conclude that the sk-strings learner is not robust.

The SMArTIC result is similar to sk-strings in that all traces generated by the simulator model X was accepted by the inferred model Y . Different from sk-strings, we noted only a slight drop in the acceptance of traces generated by Y . This drop is far less

than the noise injected ((0.1%,0.6%) vs. 4%, (0.7%,1.4%) vs. 8% and (1.9%,2.6%) vs. 10%). These indicates that filtering of erroneous traces is effective in preventing loss of precision.

The most important observation here is that: Having coring as post-processing to sk-strings removes not just erroneous transitions but also quite a fair number of correct transitions. Consequently, the accuracy of the mined specification degraded. This result strongly supports our first hypothesis.

Another limitation of coring is due to the fact that transition labels are being ignored during the coring operation. The coring method only searches for the pair of nodes (i, j) where there is a low “heat transmission” from node i to node j [6]; it ignores the detail of how the node j is reached (which can be a single transition, a set of transitions, a single path or a set of paths). In the second sub-experiment, erroneous transitions go to valid nodes instead of the special error node. This results in *little/no filtering* when coring is used.

5.3 Experiment 2 Findings

We performed two sets of scalability sub-experiments. In the first set of sub-experiments, we generated distinct models by varying the number of nodes while keeping the maximum transitions per node at 4. In the second set of sub-experiments, we varied the maximum number of transitions while keeping the total number of nodes constant at 10.

Results. The experiment results are shown below for sk-strings and SMArTIC. Columns X.N and X.TN correspond to the number of nodes and maximum number of transitions per node in the specification model.

Varying No of Nodes						
sk-strings				SMArTIC		
X.N/X.TN	Precs	Recall	PS	Precs	Recall	PS
10/4	0.437	1.000	0.560	0.584	0.988	0.690
20/4	0.003	1.000	0.015	0.185	0.998	0.227
30/4	0.0004	1.000	0.005	0.059	0.999	0.090
40/4	0.0004	1.000	0.004	0.014	1.000	0.007
Varying Max No.Of Transitions						
sk-strings				SMArTIC		
X.N/X.TN	Precs	Recall	PS	Precs	Recall	PS
10/3	0.113	1.000	0.218	0.453	0.984	0.504
10/5	0.187	1.000	0.284	0.578	0.993	0.424
10/7	0.084	1.000	0.213	0.500	0.984	0.508
10/9	0.073	0.997	0.087	0.514	0.990	0.329

Analysis The above results shows that sk-strings and SMArTIC were affected when we scaled up the model size. Comparing the two set of experiments, we observe that the precision is adversely affected in all cases when we increase the number of nodes, whereas the impact is less severe when we increase the number of transitions.

SMArTIC is generally better in terms of precision up to a factor of over 147.5 (*ie.*, 30-node case). In the second set of experiments (*ie.* when we increase the maximum number of transitions), SMArTIC maintains its precision while sk-strings loses it as the maximum number of transitions increased.

6 Related Work

Specification Miners. There have been numerous work in the area of specification mining. They can be classified into two groups, depending on how the mined specifications are represented: *automaton-based* [6, 3, 9, 5, 7] and non-automaton based [34, 8, 4, 12] specification mining.

The specification miner described in [6] has been extensively referenced in this paper. In other work, Whaley *et al.* extract object-oriented component interface sequencing constraints to form multiple finite state automatons [9]. Reiss *et al.* encode program execution traces as directed acyclic graphs to aid visualization and understanding of programs [5]. Arts *et al.* dynamically extract program models from Erlang programs as state graph models for model checking and visualization [7]. We believe that these and other similar miners can benefit from our architecture with minimal changes.

Separating Error and Normal Traces. Weimer and Necula contrasted (previously labelled) error and normal traces to statically mine and filter temporal specifications for error detection in the form of 2 event rules [38]. In our case, we use dynamic analysis to mine automaton-based specification, also unsupervised learning is used to identify and filter erroneous traces. In [39], Ammons *et al.* utilizes hierarchical filtering to aid a specification miner user to detect and delete clusters of error traces *en masse*. Their method heavily depends on user input while the method proposed in this paper is automatic.

Mining Rules for Software Engineering Tasks. In [8], association rules (*ie.*, describing association of items rather than the sequence of events) based on frequent itemset mining have been used to describe user interaction behavior for web-based GUIs.

In [34], Yang *et al.* present their work in mining a restricted form of *response* pattern [40] called an “alternating” pattern using a set of templates. They only consider rules involving two events; *ie.*, of the form $a \rightarrow b$. In order to handle longer rules, Yang *et al.* introduce “chaining”. For example, if both $A \rightarrow B$ and $B \rightarrow C$ are significant, they can be concatenated to form $A \rightarrow B \rightarrow C$ which will be significant too. However, the reverse may not be always true: $A \rightarrow B \rightarrow C$ might be significant although only rule $A \rightarrow B$ is significant while $B \rightarrow C$ is not. For such cases, the rule $A \rightarrow B \rightarrow C$ cannot be generated by inferring from two-event rules and chaining them. Hence, chaining only gives a partial solution rule set for multi-event (> 2) sequential patterns.

SMArTIC also mines a temporal *response* pattern [40] but for the purpose of generating rules for outlier detection. In our case, expressivity is required. Thus, the pre- and post-condition of our mined outliers detection rules can take in any number of events. This pattern is also referred to as a *chain response* pattern in [40]. Moreover, rather than mining all possible statistically significant rules, we only mine those useful for outlier detection.

Program Trace Characterization. Larus proposes using whole program path (WPP) to represent a program’s dynamic control flow [41]. Sequitur [33] is used to compact acyclic program trace into WPP which is its corresponding grammar. WPP is later used to find hot subpaths – heavily executed code – for optimization purpose.

Reiss *et. al* present a system for analyzing java trace data [42]. Sequitur [33] is used during path analysis to compact a trace into its grammar representation which serves as its identifier.

Similar to Larus and Reiss *et. al*, we used Sequitur to get a representation of program traces. In many applications, Sequitur is able to infer reasonable hierarchical structure from a sequence of discrete symbols. Different from Larus and Reiss *et.al*, our purpose is to obtain an intuitive measure of similarity between two traces while taking their inherent structure into consideration. Hence, rather than taking the grammar output of sequitur directly, we post-process it into a regular expression (regex) format. The regex representation “flattens” the grammar while retaining its structure. Two regex representations can then be easily compared by global sequence alignment which produces their best alignment.

Clustering Temporal Data & Similarity Measures. In clustering temporal data, there are numerous works on clustering time series data and similarity measures between two time series (*c.f.* [43]). However, there are huge difference between time series data – *e.g.* stock market data – and program trace data. The first are often lacking in structure while the later have underlying structure since they are generated from program having structured control flow.

7 Conclusion

In this paper, we began with two hypotheses about how specification miners should be organized to alleviate the impact of erroneous transitions and to localize and minimize over-generalization. We then presented our novel Specification Mining Architecture with Trace Filtering and Clustering, (SMArTIC) to support our hypotheses. SMArTIC comprises four major blocks – clustering, filtering, learning and merging. Filtering and clustering is meant to address the issue of robustness and scalability respectively.

Traces deviating from common trace population rules are removed. The resultant filtered traces are then separated into multiple clusters. By clustering common traces together, it is expected that the learner is able to learn better and over-generalization of a subset of traces is not propagated to other clusters. These clusters of filtered traces are then inputted to a specification miner. The sk-strings learner is used for learning, and each cluster is considered an independent (sub-)protocol. Lastly, a merger sub-system produces a merged automaton without sacrificing accuracy.

Along with the architecture, we have also proposed a novel trace clustering technique based on grammatical similarity, a novel outlier detection rule mining technique and a novel automaton merging method. Besides having automaton as specification, the outlier-detection rules produced by the filtering block can also be viewed as sets of simple specifications based on strong properties of significant trace group useful for filtering. They can effectively capture those property patterns proposed in [40] which are interesting for program traces and useful for identifying potential bugs.

We experimented with the Jakarta Commons Net open-source library. Our experiments aim at deriving API interaction protocols for client applications of Jakarta Commons Net open-source library [14]. From one hundred experiments performed, the following are noted: (1) SMArTIC improves precision (more than double) and co-emission while maintaining good recall (2) Both clustering and filtering help in improving precision while maintaining good recall and equivalent co-emission (3) Coring removes erroneous behavior together with a significant proportion of valid behavior – recall is reduced by more than 40% (4) Outlier detection rules are able to filter on average 43% of erroneous traces while only wrongly filter 4% of valid ones.

Further experiments using simulation measuring precision and recall in the two dimensions of increasing percentage of error (*i.e.* robustness) and increasing model size (*i.e.* scalability) of sk-strings and SMArTIC were performed. 1,800 tests on three percentage of error levels and 160 tests on different configurations of the number of nodes and the maximum number of transitions of the specification model were performed.

From the robustness experiments, the precision of sk-strings is reduced proportionally to the error induced. On the other hand, only a slight reduction of precision is observed for SMArTIC. Our experiments also show the limitation of the coring method in removing errors. From the scalability experiments, both sk-strings and SMArTIC are adversely affected by the increase in model scale (number of nodes). However, SMArTIC is able to retain better precision as compared to sk-strings up to factor of over 100.

Our experiments have strongly supported our belief that SMArTIC can produce *more precise results with good recall and equivalent, or even better, co-emission* in the presence of errors and increasingly large model.

8 Acknowledgments

We would like to thank Anand Raman, Peter Andrae and Jon D. Patrick for letting us use their sk-strings learner in our experiment. We would also like to thank Glenn Ammons and Rastislav Bodik for sharing the detail of their coring algorithm. Furthermore, we would like to thank Jianyong Wang and Jiawei Han for providing the executable of BIDE. Last but not least, we would like to thank Hugh Anderson for his help in proof-reading this paper.

References

1. Deelstra, S., Sinnema, M., Bosch, J.: Experiences in software product families: Problems and issues during product derivation. In: Proc. of Software Product Line Conf. (2004)
2. Capilla, R., Dueñas, J.: Light-weight product-lines for evolution and maintenance of web sites. In: Proc. of Euro. Conf. On Software Maintenance And Reengineering. (2003)
3. Cook, J.E., Wolf, A.L.: Discovering models of software processes from event-based data. ACM Transactions on Software Engineering and Methodology **7**(3) (1998) 215–249
4. Ernst, M., Cockrell, J., Griswold, W., Notkin, D.: Dynamically discovering likely program invariants to support program evolution. IEEE Transaction on Software Engineering **27**(2) (2001) 99–123

5. Reiss, S.P., Renieris, M.: Encoding program executions. In: Proc. of the Int. Conf. on Software Engineering. (2001)
6. Ammons, G., Bodik, R., Larus., J.R.: Mining specification. In: Proc. of Principles of Programming Languages. (2002)
7. Arts, T., Fredlund., L.: Trace analysis of Erlang program. In: Proc. of Erlang Workshop. (2002)
8. El-Ramly, M., Stroulia, E., Sorenson, P.: Interaction-pattern mining: Extracting usage scenarios from run-time behavior traces. In: Proc. of ACM-SIGKDD Int. Conf. on Knowledge Discovery and Data Mining. (2002)
9. Whaley, J., Martin, M., Lam, M.: Automatic extraction of object oriented component interfaces. In: Proc. of the Int. Symposium on Software Testing and Analysis. (2002)
10. J.Henkel, Diwan, A.: Discovering algebraic specifications from java classes. In: Proc. of the Euro. Conf. of Object Oriented Programming. (2003)
11. Fox, A.: Addressing software dependability with statistical and machine learning techniques. In: Proc. of Int. Conf. of Software Engineering. (2005) Invited Talk.
12. Raz, O., Koopman, P., M.Shaw: Semantic anomaly detection in online data sources. In: Proc. of Int. Conf. on Software Engineering. (2002)
13. Sun-Microsystems-Inc.: The java hotspot performance engine architecture. <http://java.sun.com/products/hotspot/whitepaper.html> (1999)
14. The-Apache-Software-Foundation: Jakarta commons/net. (online: <http://jakarta.apache.org/commons/net/>)
15. Nimmer, J.W., Ernst, M.D.: Automatic generation of program specifications. In: Proc. of Int. Symposium on Software Testing and Analysis. (2002)
16. Lyngsø, R., Pedersen, C., Nielsen, H.: Metrics and similarity measures for hidden Markov models. In: Proc. of National Conf. on Artificial Intelligence. (1999)
17. Law, A., Kelton, W.D.: Simulation modelling and analysis. McGraw-Hill (2000)
18. D.Lo, S-C.Khoo: QUARK: Empirical assessment of automaton-based specification miners. In: Proc. of Working Conf. on Reverse Engineering. (2006)
19. Moore, D., McCabe, G.: Introduction to the Practice of Statistics. W.H. Freeman & Co., Third Edition (1998)
20. Gupta, N.: Generating test data for dynamically discovering likely program invariants. In: Proc. of the workshop on dynamic analysis. (2003)
21. Ntafos, S.: A comparison of some structural testing strategies. IEEE Transactions on Software Engineering **14**(6) (1988) 868–874
22. van Rijsbergen, C.: Information Retrieval. Butterworths, London (1979)
23. Keogh, E., Lonardi, S., Ratanamahatana, C.: Towards parameter-free data mining. Proc. of Int. Conf. on Knowledge Discovery and Data Mining (2004)
24. Huth, M., Ryan, M.: Logic in Computer Science. Cambridge (2004)
25. Han, J., Kamber, M.: Data Mining Concepts and Techniques. Morgan Kaufmann (2001)
26. Wang, J., Han, J.: BIDE: Efficient mining of frequent closed sequences. In: Proc. of Int. Conf. on Data Engineering. (2004)
27. Yan, X., Han, J., Afhar, R.: CloSpan: Mining closed sequential patterns in large datasets. In: Proc. of SIAM Int'l Conf on Data Mining. (2003)
28. Spiliopoulou, M.: Managing interesting rules in sequence mining. In: Proc. of Euro. Conf. on Principles and Practice of Knowledge Discovery in Databases. (1999)
29. Kaufman, L., Rousseeuw, P.: Finding Groups in Data: an Introduction to Cluster Analysis. John Wiley & Sons (1990)
30. Foss, A., Zaiane, O.: A parameterless method for efficiently discovering clusters of arbitrary shape in large datasets. In: Proc. IEEE Int. Conf. on Data Mining. (2002)
31. Tompa, M.: Lecture notes on biological sequence analysis. Technical Report 2000-06-01 University of Washington (2000)

32. Cormen, T., Leiserson, C., Rivest, R., C.Stein: Introduction to Algorithms. MIT Press (2001)
33. Nevill-Manning, C., Witten, I., Maulsby, D.: Compression by induction of hierarchical grammars. In: Proc. of Data Compression Conf. (1994)
34. Yang, J., Evans, D., Bhardwaj, D., Bhat, T., M.Das: Perracotta: Mining temporal API rules from imperfect traces. In: Proc. of Int. Conf. on Software Engineering. (2006)
35. Raman, A.V., Patrick, J.D.: The sk-strings method for inferring PFSA. In: Proc. of the workshop on automata induction, grammatical inference and language acquisition. (1997)
36. Biermann, A., Feldman, J.: On the synthesis of finite-state machines from samples of their behaviour. IEEE Transactions on Computers **21** (1972) 591–597
37. JRat: Jrat the java runtime analysis toolkit. (online: <http://jrat.sourceforge.net/>)
38. W.Weimer, G.Necula: Mining temporal specifications for error detection. In: Proc. of Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems. (2005)
39. Ammons, G., Mandelin, D., Bodik, R., Larus., J.: Debugging temporal specifications with concept analysis. In: Proc. of Programming Language Design and Implementation. (2003)
40. Dwyer, M., Avrunin, G., Corbett, J.: Patterns in property specifications for finite-state verification. In: Proc. of Int. Conf. on Software Engineering. (1999)
41. Larus, J.: Whole program paths. In: Proc. of Programming Language Design and Implementation. (1999)
42. Reiss, S., Renieris, M.: Generating java trace data. In: Proc. of the ACM Conf. on Java Grande. (2000)
43. Keogh, E., Kasetty, S.: On the need for time series data mining benchmarks: A survey and empirical demonstration. Data Mining and Knowledge Discovery **7** (2003) 349–371

9 Appendix

The following figures are procedures referred to by AutomatonMerge algorithm (outlined in Figure 11 at subsection 3.4).

```

Procedure CreateUnifiableList
Inputs:
  X : First automaton
  Y : Second automaton
  td : top-down or bottom-up merge
Outputs:
  Uni: Unifiable mapping of nodes in X and Y
Method:
  Let nodesX = List of nodes by traversing X breadth first
  Let nodesY = List of nodes by traversing Y breadth first
  Let Uni = Empty If (!td) Reverse nodesX and nodesY
  Let bChange = T
  Let PrevDep = Empty
  while (bChange) do
    Let Dep = Empty
    Let PairSet = PrevDep is Empty? X.Nodes×Y.Nodes: PrevDep
    For every (nx, ny) in PairSet do
      bChange = bChange || Call CheckUnifiable (nx,ny,Uni,Dep,td)
    end for
    bChange = bChange || Call SolveCircDep (Dep, Uni)
    PrevDep = Dep
  end while
Output Uni

```

Fig. 17. Create Unifiable List

Procedure CheckUnifiable**Inputs:***nodeX* : A node in automaton *X**nodeY* : A node in automaton *Y**Uni* : Unifiable nodes*Dep* : Dependencies in CNF format*td* : top-down or bottom-up merge**Outputs:***bUnifiable*: Node Unifiable?T:F**Method:**Let *txS* = {}Let *tyS* = {}If (*!td*) Set *txS* = *nodeX*.Next Set *tyS* = *nodeY*.Next

else

 Set *txS* = *nodeX*.Prev Set *tyS* = *nodeY*.Prev

end if

Let *bUnifiable* = Fif (Unique(*txS*) == Unique(*tyS*)) if (*Uni*.Unifiable(*nodeX*,*nodeY*)) *bUnifiable* = T

end if

 Set *bUnifiable* = *bUnifiable* || Call GenCons(*txS*,*tyS*, *Uni*, *Dep*) if (LabelRepeatExist(*txS*) || LabelRepeatExist(*tyS*)) Set *bUnifiable* = *bUnifiable* && Call GenCons(*tyS*,*txS*,*Uni*,*Dep*)

end if

else *bUnifiable* = Fif (*bUnifiable*) Add (*nodeX*,*nodeY*) to *Uni***Output** *bUnifiable***Fig. 18.** Check Unifiable

Procedure GenCons**Inputs:***TransToCheckX* : *X* transitions*TransToCheckY* : *Y* transitions*Uni* : Unified list of nodes*Dep* : Dependencies in CNF form**Outputs:***bUnifiable*: Node Unifiable?T:F**Method:**

Let AND = Empty set of AND-ed constraints

for every *tx* from *TransToCheckX*

Let OR = Empty set of OR-ed constraints

 for every *ty* from *TransToCheckY* if (*tx.Lbl*==*ty.Lbl*&&!*Uni.Unifiable*(*tx.Sink*,*ty.Sink*)) OR.Add ((*tx.Sink*,*ty.Sink*))

end if

end for

AND.Add (OR)

end for

Dep.Add ((*tx.Source*,*ty.Source*),AND)Let *bUnifiable* = (#AND == 0)? T:F**Output** *bUnifiable***Fig. 19.** Generate Constraint

```

Procedure SolveCircDep
Inputs:
Dep : Dependencies in CNF format
Uni : Unifiable node list
Outputs:
bExist: Circular dep. exists?T:F
Method:
Let bExist = F
for (i=0;i<#Dep;i++)
  Let Pair = Dep.GetDep(i).Pair
  Let AND = Dep.GetDep(i).AND
  for at most #Dep times do
    if ( $\forall$  OR  $\in$  AND, (OR.Contains(Pair) ||  $\exists$  cons  $\in$  OR, Uni.Unifiable (cons)))
      bExist = T
      Uni.Add (Pair)
    end if
    Let NxtAND = Empty AND-ed dependencies
    for each OR  $\in$  AND
      if (OR.Contains(Pair))
        NxtAnd.Add (OR)
        Continue
      end if
      Let NxtOR = Empty OR-ed dependencies
      for each cons in OR
        Let NxtCons = Dep.Get (cons)
        NxtOr.Add (NxtCons)
      end for
      Let NxtOrCNF = Convert NxtOr to CNF
      NxtAND.AddRange.W/o.Duplicate(NxtOrCNF)
    end for
    Set AND = NxtAND
  end for
end for
Output bExist

```

Fig. 20. Solve Circular Dependency

Procedure CreateStructuralMerge**Inputs:**

XStart : First automaton (*X*) start node
YStart : Second automaton (*Y*) start node
MStart : Merged automaton start node
WX : Weight of automaton *X*
WY : Weight of automaton *Y*
Uni : Unifiable nodes
XUsed, *YUsed*: Nodes used for merging (so far)
TransXUsed, *TransYUsed*: Merged trans. (so far)

Method:

```
Let TransToCheckX = XStart.Next
Let TransToCheckY = YStart.Next
for each tx in TransToCheckX
  Let nX = tx.Sink
  Let bestScore = 0
  Let tBest = Null
  for each ty in TransToCheckY
    Let nY = ty.Sink
    if (Uni.Unifiable (nX, nY))
      Let PotSet = {(tx, ty) ∈ nX.Next × nY.Next | tx.Lbl == ty.Lbl ∧
                    Uni.Unifiable(tx.Sink, ty.Sink)}
      Let score = |PotSet|
      Update bestScore & tBest
    end if
  end for
end for
if (bestMatch != NULL)
  TransXUsed.Add (tx)
  TransYUsed.Add (tBest)
  Let m = XUsed.Get (tx.Sink)
  Let bRecurse = F
  if (m == NULL)
    Set m = Create new node
    XUsed.Add (tx.Sink, m)
    YUsed.Add (tBest.Sink, m)
    Set bRecurse = T
  end if
  Let dWX = WX / (WY + WX)
  Let dWY = WY / (WY + WX)
  Let prob = dWX * tx.Prob + dWY * tBest.Prob
  MStart.Next.Add (m, tx.Lbl, prob)
  if (bRecurse)
    Call CreateStructuralMerge (tx.Sink, tBest.Sink, m, WX, WY, Uni, XUsed, YUsed,
    TransXUsed, TransYUsed)
  end if
end for
```

Fig. 21. Create Structural Merge

Procedure CreateFullMerge**Inputs:***node* : A node*DSSet* : Set of Data Structure*TVS*: List of node traversed before*SHR*: Shared nodes in structural merge*W*: Weight*NUsed,TransUsed*: Merged nodes and trans. of *X*
where *node* $\in X$ **Method:**if (*node* $\in TVS$) returnelse *TVS*.Add (*node*)Let *mX* = *NUsed*.Get (*node*)Let *mW* = *mX* $\in SHR?W:1$ for each *tx* in *node*.Next doLet *nextX* = *tx*.Sinkif (!*tx* $\in TransUsed$)if (! \exists (*nextX*,*d*) $\in NUsed$)if (\exists (*nextX*,*d*) $\in DSSet$)Let *aDS* = *DSSet*.Get (*nextX*)Let *nextX* = *aDS*.ReprObj*TransUsed*.Add (*tx*)if (! $\exists t \in mX.Next | t.Sink == aDS.ReprObj \wedge tx.Lbl = t.Lbl$)*t*.AddProb (*mW* * *tx*.Prob)else *mX*.Next.Add (*nextX*,*tx*.Lbl,*mW***tx*.Prob)

end if

NUsed.AddRange (*aDS*.DescMap)*TransUsed*.AddRange (*aDS*.TransSet)

else

Let *nextRep* = Create new node*mX*.Next.Add (*nextRep*,*tx*.Lbl,*mW***tx*.Prob)*NUsed*.Add (*nextX*, *nextRep*)*TransUsed*.Add (*tx*)

end if

else

Let *nextRep* = *NUsed*.Get (*nextX*)if (! $\exists t \in mX.Next | t.Sink == nextRep \wedge tx.Lbl = t.Lbl$)*t*.AddProb (*mW* * *tx*.Prob)else *mX*.Next.Add (*nextRep*,*tx*.Lbl,*mW***tx*.Prob)

end if

end if

end if

Call CreateFullMerge (*nextX*,*DSSet*,*TVS*,*SHR*,*W*,*NUsed*,*TransUsed*)

end for

Fig. 23. Create Full Merge