

Scalable Parallelization of Specification Mining Using Distributed Computing

Shaowei Wang¹, David Lo¹, Lingxiao Jiang¹, Shahar Maoz², Aditya Budi¹

¹*Singapore Management University, Singapore*, ²*Tel Aviv University, Israel*

Abstract

Mining specifications from logs of execution traces has attracted much research effort in recent years since the mined specifications, such as program invariants, temporal rules, association patterns, or various behavioral models, may be used to improve program documentation, comprehension, and verification. At the same time, a major challenge faced by most specification mining algorithms is related to their scalability, specifically when dealing with many large execution traces.

To address this challenge, we present a general, distributed specification mining algorithm that can parallelize and distribute repetitive specification mining tasks across multiple computers to achieve speedup proportional to the number of machines used. This general algorithm is designed based on our observation that most specification mining algorithms are data and memory intensive while computationally repetitive. To validate the general algorithm, we instantiate it with five existing sequential specification mining algorithms (CLIPPER, Daikon, k-tails, LM, and Perracotta) on a particular distributed computing model—MapReduce and one of its implementations Hadoop, to create five parallelized specification mining algorithms, and demonstrate the much improved scalability of the algorithms over many large traces ranging from 41MB to 157GB collected from seven DaCapo benchmark programs. Our evaluation shows that our parallelized Perracotta running on four machines (using up to eight CPU cores in total) speeds up the original sequential one by 3 to 18 times; The other four sequential algorithms are unable to complete analyzing the large traces, while our parallelized versions can complete and gain performance improvement by utilizing more machines and cores. We believe that our general, distributed algorithm fits many specification mining algorithms well, and can be instantiated with them to gain much performance and scalability improvements.

Email address: {shaoweiwang.2010,davidlo,lxjiang,adityabudi}@smu.edu.sg, maoz@cs.tau.ac.il (Shaowei Wang¹, David Lo¹, Lingxiao Jiang¹, Shahar Maoz², Aditya Budi¹)

Keywords: Dynamic Analysis, Execution Profiles, Hadoop, MapReduce, Parallelization, Scalability, Specification Mining

Methods:

- Specification Mining Algorithms
 1. Frequent Pattern-Based Specification Miner
 2. Value-Based Invariant Miner
 3. Finite-State Machine Specification Miner
 4. Live Sequence Chart Miner
 5. Temporal Rule Miner
- Distributed Computing Models
 1. Message-Passing
 2. MapReduce
 3. Hadoop

1. Introduction

Specification mining is a family of program analysis techniques that extract likely specifications from code or execution traces. Specifications refer to certain patterns or properties that should hold in a program. They can take various forms, such as temporal rules about the order of certain method calls, and invariants that constrain method parameters and return values. The extracted specifications can provide much information about program properties which are not explicitly documented, and can be used to improve program documentation, comprehension, and verification tasks [1].

An important challenge for many specification mining algorithms relates to their scalability since they need to take many potentially large program behavioral profiles as input to search for common patterns. A common way used to collect behavioral profiles is to execute a subject program with many test cases. In order to exercise the many behaviors of a large program, many test cases would need to be run. Also, the resultant execution traces are likely to be huge. The size of a code base, the number of test cases, the sizes of generated traces are all hurdles to the scalability of existing specification mining algorithms. For example, our evaluation on four existing specification mining algorithms, (1) CLIPPER [2], a recurring pattern mining algorithm, (2) Daikon [3], a value-based invariant mining algorithm, (3) k-tails [4, 5], a finite-state machine inference algorithm, and (4) LM [6], a sequence diagram mining algorithm, shows that they fail to analyze the large traces ranging from 41MB to 157GB generated from seven DaCaPo benchmark programs [7]. A fifth algorithm, Perracotta [8], a temporal rule mining algorithm, takes hours to analyze the traces before producing some specifications. In order to

analyze many large traces from a large code base, existing specification mining algorithms need to be made much more scalable.

We observe that most specification mining algorithms are data intensive on one hand but computationally relatively repetitive on the other hand, and many repetitive tasks in those algorithms can be executed concurrently. Even though tasks in different algorithms may require various synchronization, the needed synchronization can be minimized with careful arrangements of the tasks to facilitate speedup when the algorithms are distributed onto multiple computers. This is the main insight that drives this chapter to address the scalability issue of many existing specification mining algorithms. Similar observations and ideas have been proposed to parallelize various algorithms in scientific computing, software engineering, data mining, and many other domains (e.g., [9, 10, 11, 12, 13]). However, as far as we know, there is little prior study on parallelization of various kinds of specification mining algorithms.

To help address the challenge of making various existing specification mining algorithms more scalable,¹ we propose a general specification mining algorithm that can perform repetitive specification mining tasks across multiple computers based on a general distributed computing model. The general algorithm is designed in such a way that itself abstracts away specific algorithmic details but captures the essences of many existing specification mining algorithms that mine specifications from program execution traces. We present this algorithm in the context of a message-passing based distributed computing model, in particular MapReduce. An algorithm designer can transform a *sequential* specification mining algorithm to a *distributed* one by following the guided steps in our algorithm and instantiating it with concrete algorithm-specific details.

To evaluate our general algorithm, we instantiate it with five existing sequential specification mining algorithms on top of a popular distributed computing model—MapReduce (MR) [14] and one of its open-source implementations—Hadoop [15], and evaluate the scalability of the distributed versions of these algorithms. In particular, we show how we follow the guidance of our general algorithm, and use a common input-trace splitting scheme and several algorithm-specific techniques to divide and conquer five specification mining algorithms, (1) CLIPPER [2], (2) Daikon [3], (3) k-tails [4, 5], (4) LM [6], and (5) Perracotta [8], and transform them into distributed ones. Note that the five algorithms produce different kinds of specifications expressed in different target languages, such as frequent patterns, value invariants, finite-state machines, sequence diagrams, and temporal rules. We evaluate the distributed algorithms on seven Java programs from the DaCapo benchmark [7] whose traces range from 41MB to 157GB. The results are encouraging. Perracotta’s distributed version implemented within MapReduce (Perracotta^{MR}) running on

¹Note that it is not our goal to improve the accuracy of existing specification mining algorithms in inferring correct specifications. Rather, our goal is to improve their *scalability*. When evaluating the accuracy of the parallelized algorithms, we only need to compare it against their sequential versions, instead of human developers’ ground truth.

four machines (using up to eight CPU cores in total) can speed up the original version by 3 to 18 times. The four other original algorithms are unable to analyze the large traces, while their distributed versions (CLIPPER^{MR}, Daikon^{MR}, k-tails^{MR}, and LM^{MR}) can complete within hours, and gain more performance improvement when more machines are employed.

Our main finding is that many specification mining algorithms fit distributed computing models well as they are comprised of many repetitive computational tasks dealing with data that may be split into partitions with limited overlapping. Our general algorithm also captures the essence of many specification mining algorithms well and can be used to help transform sequential algorithms into distributed ones to gain much performance and scalability improvements by implementing them within the MapReduce framework and executing them on clusters of computers. We believe our findings are applicable to many other specification mining algorithms, especially those that mine specifications expressed in one of the five target languages that we have investigated.

The contributions of this chapter are as follows:

1. Similar to many prior studies on parallelization of other algorithms in various domains, we observe that many *specification mining* algorithms can be fit into a distributed programming model, and much performance and scalability gains can be achieved by parallelizing them within a distributed computing framework, such as MapReduce.
2. We present a general distributed specification mining algorithm that abstracts away particular algorithmic details and represents the essences of many existing specification mining algorithms.
3. We propose an input-trace splitting scheme and several algorithm-specific techniques to instantiate the general algorithm with five existing sequential specification mining algorithms to create five distributed algorithms.
4. We perform an empirical evaluation with seven Java programs from the DaCapo benchmark and show that the five distributed algorithms perform significantly faster than the original algorithms on many large traces.

This chapter is organized as follows. Section 2 provides a brief introduction to the five specification mining approaches and the distributed computing model we use in our work. Section 3 presents the main technical contribution of the chapter, that is, the general distributed specification mining algorithm and its instantiations with the five existing algorithms. Our implementation and empirical evaluation are described in Section 4. Section 5 discusses related work. Section 6 concludes with future work.

2. Background

In this section, we first briefly introduce each of the five mining algorithms that we parallelize. Then, we introduce the distributed computing model used in this chapter—the message-passing model and MapReduce.

2.1. Specification Mining Algorithms

Based on the format of the specifications that a specification mining algorithm produces [1], many algorithms can be grouped into ones that produce (1) frequent patterns, (2) value-based invariants, (3) finite-state machines, (4) sequence diagrams, and (5) temporal rules. We briefly describe these families of algorithms in the following.

We present sample outputs of the five kinds specification mining algorithms in Figure 1, which can be mined from various kinds of program execution traces.

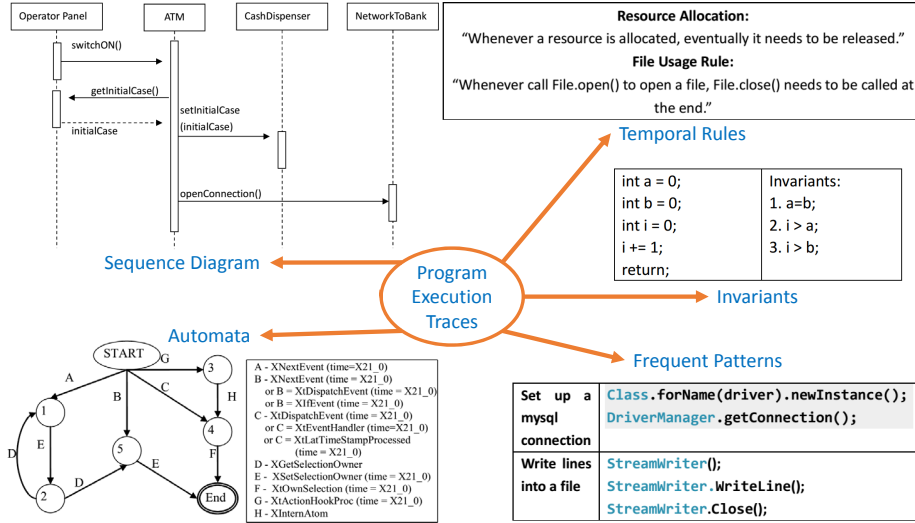


Figure 1: Sample Outputs of Specification Mining Algorithms

2.1.1. Mining Frequent Patterns

Discovering patterns that appear many times in large input datasets is a well-known problem in data mining [16]. Many algorithms, such as frequent itemset mining, sequential pattern mining, and graph pattern mining, aim to capture frequent patterns. A number of algorithms specific to software engineering tasks have been proposed. For example, interaction pattern mining [17] analyzes traces of system-user interactions to discover frequently recurring activities and uses them as parts of functional requirements for re-engineering. Iterative pattern mining (CLIPPER) [2] takes in a set of execution profiles containing methods invoked during the executions and then identifies methods that often need to be invoked together or in a particular order as usage specifications for the methods.

2.1.2. Mining Value-Based Invariants

A value-based invariant captures the relation (e.g., $x==y$) among program variables that should be satisfied at a program point (e.g., when a method

returns). Daikon is the pioneer and most well-known work extracting value-based invariants [3]. It has many invariant templates, such as Equality (e.g., `x==y`), `IntGreaterThan` (e.g., `iVal1>=iVal2`), `IntArraySorted` (e.g., `isSorted(iArray1)`), etc. Based on a set of input execution traces, Daikon matches the traces to the templates at various program points of interests (e.g., method entries and exits). Instances of the invariant templates satisfied by all (or most) of the input traces are outputted.

Value-based invariants generated by Daikon can be used independently, or be used in conjunction with other kinds of specifications, e.g., to enrich finite-state machines [5] or sequence diagrams [18].

2.1.3. Mining Finite-State Machines

Many of these algorithms extend or make use of techniques from the grammar inference community [19, 20, 5]. One of these algorithms, `k-tails` [20], builds a prefix tree acceptor from a set of execution traces that capture input-output behaviors; the nodes of the prefix tree acceptor are then merged based on some evaluation criteria, e.g., the similarity of the subsequent k -paths whose lengths are at most k , to form finite-state machines, which are then used as specifications of program behaviors.

2.1.4. Mining Sequence Diagrams

Sequence diagrams are a visual formalism to specify the ordering of events among components in a system. Different algorithms exist for mining various kinds of sequence diagrams, such as UML sequence diagrams [21], message sequence charts [22], message sequence graphs [4], live sequence charts (LSCs) [23, 24, 6], etc. Such visual diagrams can help maintainers of a program to better understand how various components in the program interact with each other.

2.1.5. Mining Temporal Rules

Temporal rules can be expressed in various formats, such as association rules [25, 8], temporal logics [26, 27], “Whenever x_1, \dots, x_n occur, y_1, \dots, y_m also occur”, etc. Such rules help to make it clearer which operations should or should not occur in certain orders so that maintainers of the program may make changes accordingly. Most temporal rule mining algorithms evaluate the validity of a rule based on the likelihood that the x s are followed by the y s, and the number of times x followed by y in the execution traces. They mainly differ in the semantics of the mined rules, the allowable values of n and m , and the metrics used to evaluate rule validity. For example, Perracotta extracts association rules of short length (n and m being 1) [8]; Others extract temporal rules of longer lengths [27].

2.2. Distributed Computing

Similar to many prior studies on parallelization of other algorithms in various domains, we observe that many specification mining algorithms can be

broken into computational tasks that are repetitively applied to various parts of input data, and thus fit well to a distributed computing model. This section summarizes the concepts we need.

2.2.1. Message-Passing Model

We focus on distributed algorithms in the message-passing model where multiple processes on multiple computing nodes have their own local memory and communicate with each other by message passing, although our general algorithm may be adapted to other distributed computing models as well.

The processes share information by *sending/receiving* (or *dispatching/collecting*) data to/from each other. The processes most likely run the same programs, and the whole system should work correctly regardless of the messaging relations among the processes or the structure of the network. A popular standard and message-passing system is Message Passing Interface (MPI) defined in [13]. Such models themselves do not impose particular restrictions on the mechanism for messaging, and thus gives programmers much flexibility in algorithm/system designs. However, this also means that programmers need to deal with actual sending/receiving of messages, failure recovery, managing running processes, etc.

2.2.2. MapReduce

MapReduce is a simplified distributed computing model for processing large data in a cluster of computers [14], reducing programmers' burden of dealing with actual sending/receiving of messages and various system issues so that programmers may focus more on the algorithmic issues. It can be implemented on top of a message-passing system, such as MPI [28]. In this chapter, we base our implementation on Hadoop [15], a free and open-source implementation of MapReduce.

The model splits the problem at hand into smaller sub-problems as requested, distributes these sub-problems among the computers in the cluster, and collects and combines the results that are passed back. Besides the splitting function, the key to use the MapReduce framework is to define two functions: (1) *map*, which takes one input key/value pair (K_{ip}, V_{ip}) , and generates zero or more intermediate key/value pairs $(\text{list}(K_{int}, V_{int}))$, and (2) *reduce*, which composes all intermediate values associated with the same key into a final output. The splitting function, customizable to split input data differently, partitions the whole input dataset into small pieces, and transforms each piece into a set of key/value pairs.

MapReduce works by automatically applying the *map* function on each of the key/value pairs from the splitting function to produce an intermediate set of key/value pairs. It then automatically groups all intermediate values associated with the same key together; the *reduce* function is then applied to each group to result in a partial output O_{part} ; all partial outputs are concatenated to form the final output. The inputs and outputs of the *map* and *reduce* functions are illustrated in Table 1. The following sections use the same symbols to represent the inputs and outputs of the functions as well.

Table 1: Map and Reduce Operations

Operation	Input	Output
<i>map</i>	(K_{ip}, V_{ip})	$\text{list}(K_{int}, V_{int})$
<i>reduce</i>	$(K_{int}, \text{list}(V_{int}))$	O_{part}

3. Distributed Specification Mining

We now present the main contribution of this chapter: a general distributed specification mining algorithm and redefinitions of five concrete algorithms in the MapReduce model.

3.1. Principles

We first present our general principles and algorithm on parallelizing specification mining algorithms.

3.1.1. Abstracting Specification Mining Algorithms

Even though many specification mining algorithms are not initially designed to be distributed, we can *divide-and-conquer* by exploiting the parallelism in various parts of the algorithms based on our observation that many computational tasks in the algorithms are repetitively applied to various parts of input data. Fig. 2 illustrates the design idea of our general, distributed specification algorithm.

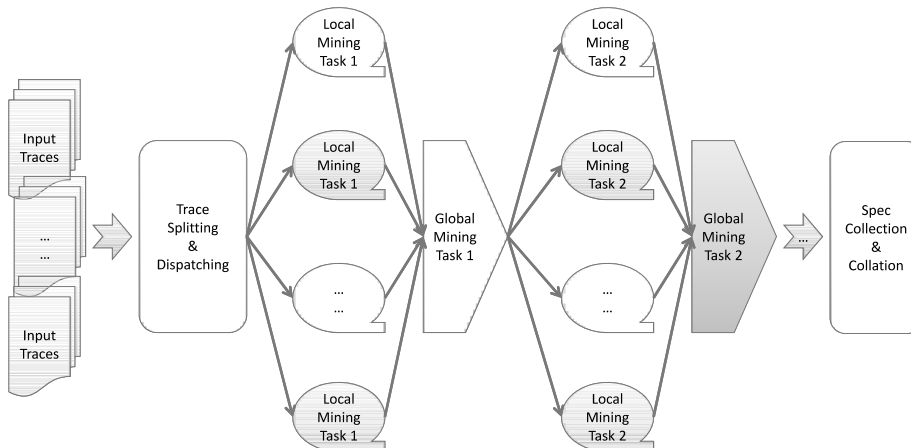


Figure 2: Overview of Our Distributed Specification Mining Algorithm

The key is to extract such tasks from existing algorithms that are repetitively applied to different parts of the input traces so that the input traces can be split and dispatched to and processed at different computing nodes. We note that many algorithms contain *local mining tasks* that can be done completely on a small part of the input data without the need of other parts. For some

algorithms, however, there are still *global mining tasks* that need to operate on all data and we need to ensure those tasks can run scalably. Fortunately, we note that the mining algorithms rely on various “statistics” that measure the likelihood of a candidate specification to be valid. It is rarely necessary for the mining algorithms to really operate on all data at once in memory. Thus, we may also split the data (either the input traces or intermediate results from other tasks) needed by the global mining tasks so that they operate on smaller data and become more parallelizable and scalable; or, we can replace the global mining tasks with certain local ones plus certain specification compositions since many specifications are compositional. Multiple iterations of local and global mining tasks may be interleaved to find specifications minded by the normal sequential algorithms.

The general steps of our approach for parallelizing a given specification mining algorithm that takes a set of execution traces as input are as follows:

- (1) Extract “local” operations in the algorithm that can be done in a separate trunk of the traces. The boundaries of trace trunks can be defined based on the operations and be algorithm-specific.
- (2) Extract “global” operations in the algorithm that may need to be done with all data and decide how the data needed by the global operations may be split and/or replace the global operations with local ones.
- (3) Split the input traces into trunks accordingly, and dispatch them to different computing nodes for either local or global operations.
- (4) Collect results from different computing nodes, and compose them to produce final specification outputs.

To produce efficient distributed versions of the sequential specification mining algorithms, one needs to ensure that the extracted local/global operations can be independent and executed concurrently with little or no synchronization. The steps described here are generic, although many details (what the local and global operations are, how to split data, how to dispatch/collect, how to compose results, etc.) are algorithm-specific, which are further explained in Section 3.2.

3.1.2. Distributed Specification Mining With MapReduce

MapReduce simplifies the general distributed computing model by providing automated mechanisms for setting up a “master” process that manages work allocated to “worker” processes, dispatching work to a worker, collecting results from a worker, recovering from failures, utilizing data locality, etc. We further describe our general specification mining steps in the context of MapReduce as follows:

- (1) Define an appropriate *map* function that corresponds to a local mining task. Each instance of the *map* function runs in parallel with respect to other instances; it takes one trace trunk V_{ip} as input to produce a set of intermediate specification mining results (*intermediate key/value pairs*,

$\text{list}(K_{int}, V_{int})$, in MapReduce’s terminology). The map function must be designed in such a way that the operation on V_{ip} is independent on other trace trunks.

- (2) Define an appropriate *reduce* function that corresponds to a global mining task or a composition operation that combines results (i.e., the intermediate key/value pairs, $\text{list}(K_{int}, V_{int})$) from local mining tasks. We note that many algorithms rarely need global mining tasks and the composition operations may be as simple as concatenation or filtering or recursive applications of some local mining tasks (see algorithm-specific steps in Section 3.2).
- (3) Define an appropriate *record reader* that splits input traces into trunks suitable for the map function. For example, if the map function from a mining algorithm deals with invariants within a method, a trace may be split at method entry and exit points. Each trace trunk can be identified by a *trace identifier* K_{ip} and its content V_{ip} .
- (4) Let the MapReduce framework automatically handles the actual trace splitting, dispatching, and result collection.

We note that the above general steps provide guidance to make it easier to transform sequential specification mining algorithms into distributed ones, although strategies and techniques used to identify the local/global tasks in various algorithms might be different from each other, and there can be multiple ways to define the local/global operations, split the data, etc. for a given algorithm.

In the following, we describe our concrete instantiations of the general algorithm on five specification mining algorithms: (1) CLIPPER [2], a recurring pattern mining algorithm, (2) Daikon [3], a value-based invariant mining algorithm, (3) k-tails [4, 5], a finite-state machine inference algorithm, and (4) LM [6], a sequence diagram mining algorithm, and (5) Perracotta [8], a temporal rule mining algorithm. We believe our findings can be easily adapted to other specification mining algorithms, especially those that mine specifications in languages similar to one of the five algorithms.

3.2. Algorithm-Specific Parallelization

3.2.1. Iterative Pattern Mining with MapReduce

We illustrate how to instantiate our general algorithm with CLIPPER, an iterative pattern mining algorithm [2], to create the distributed version CLIPPER^{MR} .

CLIPPER, similar to many frequent pattern/sequence mining algorithms, explores the search space of all possible patterns. It starts with small patterns and then *grows* these patterns to larger ones. *Pattern growth* is performed repeatedly; each iteration grows a pattern by one unit. The iterations follow the depth-first search procedure. During the traversal of search space, every pattern that is *frequent* (i.e., appearing many times in the data set) is outputted. There have been studies on parallelization of frequent sequence mining algorithms,

Algorithm 1 Generic Algorithm of Frequent Pattern Mining

```
1: Procedure MinePatterns:
2:   Let SMALL = Small frequent patterns
3:   for all s in SMALL do
4:     TraverseSSpace(s)
5:   end for
6:
7:   Procedure TraverseSSpace(Pattern s):
8:     Output s
9:     Let NEXT= GrowPattern(s)
10:    for all n in NEXT do
11:      TraverseSSpace(n)
12:    end for
13:
14:   Procedure GrowPattern(Pattern s):
15:     Let BIGGER = s++e , where e is a growth unit and ++ is a grow operation
16:     for all s' in BIGGER do
17:       if s' is frequent then
18:         Output s'
19:       end if
20:     end for
```

such as [9]. Their work uses MapReduce too, but our data sources and the subject algorithms are specific to specification mining, which requires different parallelization strategies and techniques. The semantics of the sequential patterns mined by their approaches is also different from that of iterative patterns mined by CLIPPER. Their approach relies on w -equivalency property which may hold only for their sequential patterns.

A piece of pseudocode for CLIPPER, as well as many frequent pattern mining algorithms, is shown in Algorithm 1. Intuitively, checking if a pattern is frequent or not could potentially be a parallelizable task. Unfortunately, it is not straightforward to break the pattern mining problem into *independent* tasks. On one hand, as we grow a pattern one unit at a time, if the pattern is not frequent, the longer pattern is not frequent either. In other words, some tasks can be omitted after the evaluation of other tasks and thus dependant on others. On the other hand, without the strategy of omitting longer patterns, the number of tasks grows exponentially with respect to the length of the traces.

Fortunately, we identify a common operation that is shared by these tasks, namely *pattern growth* (i.e., procedure **GrowPattern** in Algorithm 1). As pattern growth is performed many times, it is the critical operation that the mining algorithm spends much resource on. Thus, rather than trying to parallelize the whole pattern mining algorithm, we parallelize the pattern growth procedure.

The pattern growth procedure considers a pattern P and tries to extend it to patterns $P++e$ where e is a growth unit and $++$ is a growth operation (e.g., appending an event to an iterative pattern—from $\langle m1 \rangle$ to $\langle m1, m2 \rangle$).

For an iterative pattern P and trace T , we store the indices pointing to the

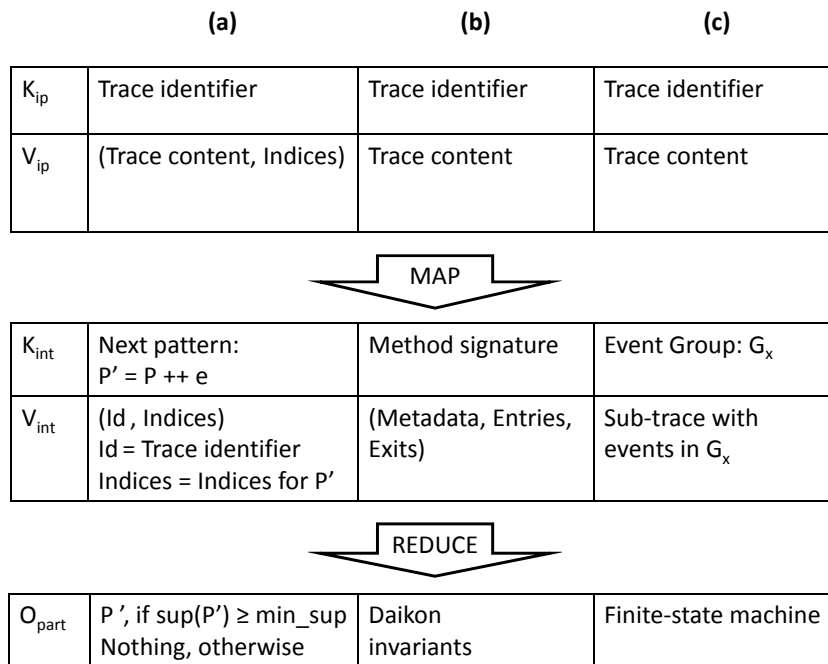


Figure 3: MapReduce inputs (K_{ip}, V_{ip}) , intermediate key/value pairs (K_{int}, V_{int}) , and outputs (O_{part}) for $\text{GrowPattern}(\text{Pattern } P)$ of CLIPPER (Column (a)), Daikon (Column (b)), and k-tails (Column (c))

various instances of the pattern in T . When we try to grow the pattern P to each pattern $P' \in \{P++e\}$, we can update these indices to point to instances of the pattern P' . From the instances of P' , we could then know if P' is frequent and thus should be in the output. Thus, we break this operation of checking *all* $P' \in \{P++e\}$ into parallelizable tasks, each of which is in the following format: check if *one* pattern P' is frequent.

We realize $\text{GrowPattern}(\text{Pattern } P)$ by instantiating the *map* and *reduce* functions in our general algorithm as follows. The *map* function works in parallel on each trace and updates the indices pointing to instances of P to indices of instances of $P' \in \{P++e\}$. It creates an intermediate key/value pair (K_{int}/V_{int}) where the key corresponds to a P' and the value corresponds to the indices pointing to instances of P' in the trace. MapReduce groups all indices corresponding to a P' . Each intermediate key K_{int} and all of its corresponding intermediate values form a task that is sent to the *reduce* function. The *reduce* function computes the support of a pattern P' and outputs it if the support is more than the minimum support threshold min_sup (i.e., if P' is frequent). We list the inputs (K_{ip}, V_{ip}) , intermediate key/value pairs (K_{int}, V_{int}) , and outputs (O_{part}) in the Column (a) in Fig. 3 for CLIPPER^{MR} .

Given a large execution trace, the pattern growth operation can be performed

in parallel. Each trace is processed in parallel by multiple instances of the map function. Also, the process to check if a pattern P' is frequent or not could be done in parallel by multiple instances of the reduce function.

Note that we only parallelize the `GrowPattern` operation, and thus each MapReduce procedure in our implementation only performs one unit of pattern growth operation (i.e., $P \rightarrow P++e$). Since many software properties are short and may be specified with only a few operation units (e.g., rules used in Static Driver Verifier [29]), we restrict the size of the patterns mined to be at most 3 to limit the experimental costs.

Example: Given two traces $trace_1 = \langle a, b, c \rangle$ and $trace_2 = \langle a, b \rangle$, we want to mine patterns with support values above $min_sup = 2$ using $CLIPPER^{MR}$. In the first iteration, $CLIPPER^{MR}$ mines patterns of length 1. We have two instances of the map function map_1 and map_2 that take as input $trace_1$ and $trace_2$, respectively. Then map_1 creates the following intermediate key/value pairs: $\{K_{int} = \langle a \rangle, V_{int} = (trace_1, \{1\})\}$, $\{K_{int} = \langle b \rangle, V_{int} = (trace_1, \{2\})\}$, and $\{K_{int} = \langle c \rangle, V_{int} = (trace_1, \{3\})\}$. Map function map_2 produces $\{K_{int} = \langle a \rangle, V_{int} = (trace_2, \{1\})\}$, and $\{K_{int} = \langle b \rangle, V_{int} = (trace_2, \{2\})\}$. We have three instances of the reduce function: $reduce_1$, $reduce_2$, and $reduce_3$. $reduce_1$ takes pairs with the key $\langle a \rangle$ and checks whether the number of instances is larger than min_sup . Similarly, $reduce_2$ and $reduce_3$ collects pairs with the keys $\langle b \rangle$ and $\langle c \rangle$, respectively. The reduce functions output patterns $\langle a \rangle$ and $\langle b \rangle$ because they satisfy the threshold min_sup . In the next iteration, $CLIPPER^{MR}$ mines patterns of length 2. The map functions generate the following intermediate key/value pairs: $\{K_{int} = \langle a, b \rangle, V_{int} = (trace_1, \{1\})\}$, $\{K_{int} = \langle a, c \rangle, V_{int} = (trace_1, \{1\})\}$, $\{K_{int} = \langle b, c \rangle, V_{int} = (trace_1, \{2\})\}$, and $\{K_{int} = \langle a, b \rangle, V_{int} = (trace_2, \{1\})\}$. The reduce functions group the instances based on the key values and find that pattern $\langle a, b \rangle$ satisfies the min_sup threshold. Finally, $CLIPPER^{MR}$ would return the following frequent patterns: $\langle a \rangle$, $\langle b \rangle$, and $\langle a, b \rangle$.

3.2.2. Value-Based Invariants Mining with MapReduce

We parallelize Daikon into a distributed version $Daikon^{MR}$ by instantiating our general algorithm in MapReduce.

Similar to Daikon, $Daikon^{MR}$ takes as input a set of execution traces and outputs invariants for each method that hold for all execution traces. We parallelize Daikon by splitting the input traces: Rather than feeding the whole set of traces to one instance of Daikon, we process the traces for each method separately, and in parallel—the trace logs for each method are fed into one instance of Daikon. This allows us to instantiate our general algorithm for Daikon relatively easily without the need for synchronization because the traces of different methods are independent from one another for inferring method-level invariants.

In $Daikon^{MR}$, the *map* function processes a set of traces and outputs $\langle method\ signature, (metadata, entries, and\ exits) \rangle$ pairs. The latter part of each pair contains method metadata (e.g., the number of parameters a method has, the types of the parameters, etc.), and parts of the execution traces corresponding to the states of the various variables when entries and exits

of the methods are executed. The *reduce* function runs an instance of Daikon on (*metadata, entries, and exits*) of the same method and outputs $\langle \text{method signature, method invariants} \rangle$ pair. We illustrate the inputs, intermediate key/value pairs, and outputs for Daikon^{MR} in MapReduce’s terminology in the Column (b) in Fig. 3.

Many instances of Daikon are executed in parallel, each of which runs on a rather small input. Thus, each instance of Daikon requires much less memory and is able to quickly produce a subset of the results.

3.2.3. Finite-State Machine Mining with MapReduce

Many finite-state machine mining algorithms are variants of the k-tails algorithm [20]. The algorithms investigate a set of execution traces and produce a single finite-state machine. However, this finite-state machine may be too large and difficult to comprehend. Many studies propose methods to *split* the traces and learn a finite-state machine for each sub-trace whose ideas can be instantiated in MapReduce with our general algorithm, and we name it as k-tails^{MR}.

Consider a mapping function: EVENTS→GROUP, where EVENTS is the set of events (i.e., method calls) in the execution traces, and GROUP is a group identifier. Events in the same group are *related*. Many notions of relatedness may be defined (c.f., [30]). In this chapter we consider one such notion: a group of events is composed of invocations of methods appearing in the same class.

The *map* function slices each trace into a set of sub-traces based on the group membership of each event. MapReduce collects the sub-traces belonging to the same group. The *reduce* function produces one finite-state machine for a group of sub-traces by invoking an instance of the k-tails algorithm. We illustrate the inputs, intermediate key/value pairs, and outputs for k-tails^{MR} in the Column (c) in Fig. 3.

The slicing could be done in parallel for separate execution traces. In addition, the learning of multiple finite-state machines could be done in parallel.

3.2.4. Sequence Diagram Mining with MapReduce

We illustrate how to transform the LM algorithm [6] that mines live sequence charts (LSC) [23, 24] into LM^{MR}.

An LSC contains two parts: pre-chart and main-chart. The semantics of LSC dictates that whenever the pre-chart is observed, eventually the main-chart would also be observed. The goal of the mining task is to find all LSCs that appear frequently (i.e., the support of the LSC is more than *min_sup*), and that the proportion of pre-chart being followed by main-chart in the execution traces is greater than a certain *min_conf* threshold. LSCs obeying these criteria are considered significant.

The LSC mining algorithm works in two steps:

- (1) Mine frequent charts;
- (2) Compose frequent charts into significant LSCs.

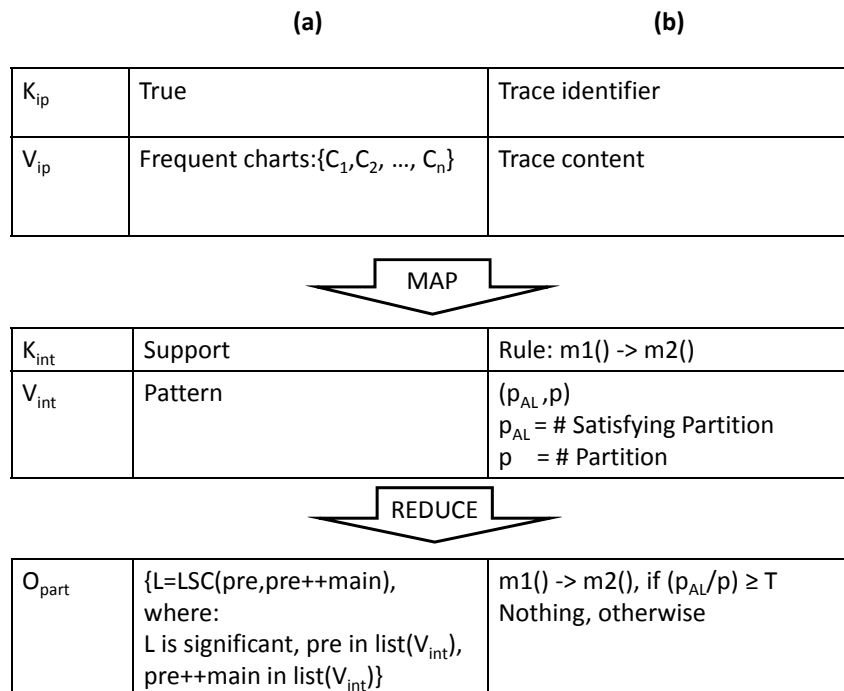


Figure 4: MapReduce inputs (K_{ip}, V_{ip}) , intermediate key/value pairs (K_{int}, V_{int}) , and outputs (O_{part}) for LM (Column (a)), and Perracotta (Column (b))

For the first step, we employ the same strategy as described in Section 3.2.1. For the second step, we consider the special case of $min_conf=100\%$ (i.e., mining of LSCs that are always observed in the execution traces—the pre-chart is always eventually followed by the main-chart).

In LSC mining, Lo and Maoz define *positive witnesses* of a chart C , denoted by $pos(C)$, as the number of trace segments that obey chart C [31]. They also define *weak negative witnesses* of C , denoted by $w_neg(C)$, as the number of trace segments that do not obey C due to end of trace being reached. The *support* of an LSC $L = pre \rightarrow main$ is simply the number of positive witnesses of $pre++main$. The *confidence* of an LSC L is given by:

$$conf(L) = \frac{|pos(pre++main)| + |w_neg(pre++main)|}{|pos(pre)|}$$

We first note that LSCs with 100% confidence must be composed of pre-chart and main-chart where $|pos(pre++main)| + |w_neg(pre++main)|$ equals $|pos(pre)|$. We could break the task of constructing all significant LSCs to sub-tasks: find all significant LSCs of a particular support value.

We name the MapReduce version of LM as LM^{MR} . For LM^{MR} , we use the following *map* and *reduce* functions. The *map* function works on the set of patterns and simply group pattern C , where either $pos(C)+w_neg(C)$ or $pos(C)$ has a particular value into a bucket. If a pattern C has different $pos(C)+w_neg(C)$ and $pos(C)$ values, it is put into two buckets. The *reduce* function constructs significant LSCs by composing two patterns in each bucket. We list the involved inputs, intermediate key/value pairs, and outputs in the Column (a) in Fig. 4.

The composition of charts into LSCs is done in parallel for separate buckets. If the significant LSCs have many different support values, the speed up in the second stage of LSC mining due to the parallelization could be substantial.

Finally, note that LM^{MR} applies MapReduce framework twice, sequentially composed in a pipeline. The first application computes the frequent charts using the solution presented in Section 3.2.1. The output of this application is used as an input for the second application described above. We note that composition of instances of MapReduce in a pipeline is also common (e.g., [32]).

3.2.5. Temporal Rule Mining with MapReduce

We illustrate how to re-implement basic Perracotta [8], by using our general algorithm and MapReduce. Perracotta proposes several extensions and variants to the main algorithm, e.g., chaining, etc. We consider only the basic Perracotta, which computes alternating properties. We call the resultant algorithm $Perracotta^{MR}$.

For n unique methods in the execution traces, Perracotta checks n^2 possible temporal specifications of the format: “Whenever method m_1 is executed, eventually method m_2 is executed” (denoted as $m_1 \rightarrow m_2$), to see if the specification is strongly observed in the execution traces. A measure known as *satisfaction rate* is defined based on the proportion of partitions in the traces that satisfy $m_1^+m_2^+$ (i.e., p) that also satisfy the temporal rule $m_1 \rightarrow m_2$ (i.e., p_{AL}). It is often the case that n is large and Perracotta would take a lot of memory to process the traces together. We break the original task into small subtasks by splitting each long trace into smaller sub-traces of size k each and process them independently—by default we set k to be 300,000 events for $Perracotta^{MR}$. As k is relatively large, by the principle of locality (i.e., related events appear close together; c.f., [33]), there will be no or little loss in the mined specifications.

Following our general algorithm, we define the following *map* and *reduce* functions. The *map* function is applied to each execution subtrace independently. For each execution subtrace, the *map* function computes for each potential rule $m_i \rightarrow m_j$ two numbers: the number of partitions in the subtrace (i.e., p) and the number of partitions in the subtrace that satisfy the rule (i.e., p_{AL}). The method pair is the intermediate key, while the two numbers p and p_{AL} are the value in the intermediate key/value pair (K_{int} and V_{int}). MapReduce groups the counts for the same rule together. The *reduce* function simply sums up the p and p_{AL} for the separate execution sub-traces and computes a *satisfaction rate* for the corresponding rule. Rules that satisfy a user-defined threshold of satisfaction

rate (i.e., S), are provided as output. By default, the satisfaction rate is 0.8. We list the involved inputs, intermediate key/value pairs, and outputs in the Column (b) in Fig. 4.

Notice that the sub-traces can now be processed in parallel using multiple runs of the *map* and *reduce* functions on potentially different machines. Also, the computation and checking of the satisfaction rate could be done in parallel. No synchronization is needed among different sub-traces.

4. Implementation & Empirical Evaluation

We have implemented the algorithms described in the previous sections in Hadoop [15], one of the most popular MapReduce implementations. We describe our datasets, experimental settings, research questions, and experimental results in the following.

4.1. Dataset & Experimental Settings

We use seven programs, *avroa*, *batik*, *fop*, *luindex*, *lusearch*, *xalan*, and *tomcat*, from the DaCapo benchmark [7] as our subjects. We also have implemented a Java instrumentation tool to collect all methods that get executed (referred to as *trace databases* later) for the experiments with CLIPPER, *k-tails*, *LM*, and *Perracotta*; we use *Chicory*, a part of *Daikon*, to collect the traces for *Daikon*. The sizes for *Daikon* traces for these seven programs range from 18GB to 157GB, while the sizes of the trace databases range from 41MB to 533MB. The experiments are run on four Acer M680G machines, each having an Intel Core i5 4-core CPU, 4GB of memory, and 2TB of hard disk, installed with the operating system Ubuntu version 12.04. One of the machines is used as the master; three others are slaves. We also configure Hadoop (version 2.0.0-alpha) to use *up to* three cores for distributed map and reduce tasks on each slave machine to reduce the effects of potential resource contentions. We set the maximum memory of each map/reduce task to 1,200 MB and leave the other settings to their default values (e.g., the Hadoop file system’s replication factor). Before running the MR versions of the specification mining algorithms, we also copy all traces from the usual ext4 file system under Ubuntu into the Hadoop file system as a one-time cost. To reduce experimental bias, we run each experiment with each version of the various specification mining algorithms two times and report the averages across the two runs.

4.2. Research Questions & Results

Our study aims to answer the following research questions:

- RQ1 Could existing specification mining algorithms scale to process large execution traces?
- RQ2 Could MapReduce be used to improve the scalability of existing specification mining algorithms?




Figure 5: Performance Improvements for CLIPPER

RQ3 How much more scalable our mining algorithms would be when we increase the number of processing cores?

We discuss the answers to the above research questions for each of the five specification mining algorithms.

4.2.1. Mining Frequent Patterns

To answer RQ1, we run the original version of CLIPPER on the traces. This version mines patterns recursively and needs to load the complete trace database into memory. Thus even for the smallest trace database with size of 41MB (from batik), original CLIPPER is unable to run.

To answer RQ2, we examine the performance of CLIPPER^{MR} with up to eight parallel map and reduce tasks. $\text{CLIPPER}^{MR(8)}$ (that is, with eight parallel map and reduce tasks) outputs the invariants for all traces from the seven programs within 1493 minutes. This shows CLIPPER^{MR} improves the scalability of the original CLIPPER.

To answer RQ3, we compare the time cost for CLIPPER^{MR} as we increase the number of parallel tasks (see Fig. 5). We find that the performance improves as we increase the number of parallel tasks. By increasing the number of parallel MapReduce tasks from one to four, we gain a speed up in range of 1.4 to 3.2 times. By increasing the number of parallel MapReduce tasks from one to eight, we gain a speed up in range of 1.7 to 4.6 times. The reason why CLIPPER^{MR} cannot speed up as much as parallelized versions of the other mining algorithms (see below) is that it needs to process a lot of I/O operations across different nodes in Hadoop system.

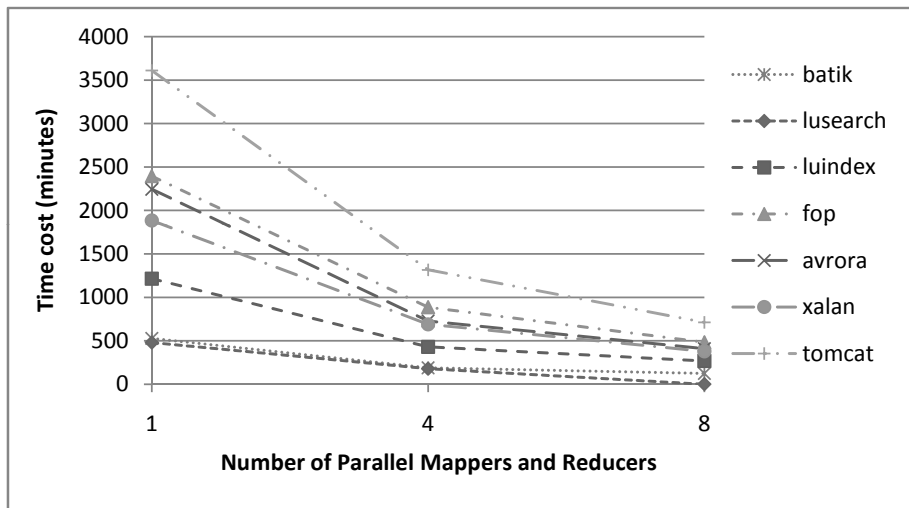


Figure 6: Performance Improvements for Daikon

4.2.2. Mining Value-Based Invariants

To answer RQ1, we run the original Daikon on the traces. Since the traces from the seven programs are all larger than 18GB, the original Daikon runs out of memory before outputting any invariant.

To answer RQ2, we examine the performance of the original Daikon with that of Daikon^{MR} with up to eight parallel map and reduce tasks. $\text{Daikon}^{MR(8)}$ outputs the invariants for *all* traces from seven programs within 2374 minutes, and we are only unable to infer the invariants of less than 5% of the methods (since we terminate a Daikon instance if it takes more than 600 seconds to complete). Obviously, Daikon^{MR} improves the scalability of the original Daikon.

To answer RQ3, we compare the time cost for Daikon^{MR} as we increase the number of parallel tasks (see Fig. 6). We find that the performance improves when we increase the number of parallel tasks. By increasing the number of parallel MapReduce tasks from one to four, we gain a speed up in range of 2.7 to 3 times. By increasing the number of parallel MapReduce tasks from one to eight, we gain a speed up in range of 4.2 to 5.4 times. We notice that the rate of speed up decreases as we increase the number of parallel tasks from 4 to 8. This is so as there are more resource contentions between the mappers and reducers as the number of parallel tasks are increased in our small four-machine cluster with limited memory.

4.2.3. Mining Finite-State Machines

To answer RQ1 and RQ2, we compare the performance of the original k-tails with that of $\text{k-tails}^{MR(8)}$. The original k-tails ran out of memory before outputting any finite-state machines. On the other hand, $\text{k-tails}^{MR(8)}$ is able to output finite-state machines for *all* programs in 40 minutes. Similar to




Figure 7: Performance Improvements for k-tails

Daikon^{MR}, we also employ a timeout and terminate an instance of k-tails construction process run in one reducer if it does not complete within 120 seconds. We find that we are unable to run k-tails to completion for only 5% of the classes. Obviously, k-tails^{MR} improves the scalability of the original k-tails.

To answer RQ3, we compare the time cost for k-tails^{MR} as we increase the number of parallel tasks (see Fig. 7). We find that the performance improves as we increase the number of parallel tasks. By increasing the number of parallel MapReduce tasks from one to four, we gain a speed up in range of 2 to 3.7 times. By increasing the number of parallel MapReduce tasks from one to eight, we gain a speed up in range of 2.3 to 5.6 times.

4.2.4. Mining Sequence Diagrams

To answer RQ1 and RQ2, we compare the performance of the original LM with that of LM^{MR(8)}. Original LM is unable to run because of memory problems, while LM^{MR} is able to get the sequence diagrams (since LM^{MR} is based on CLIPPER^{MR}). LM^{MR(8)} can output the invariants for all traces from the seven programs within 1508 minutes. This shows LM^{MR} can improve the scalability of original LM.

To answer RQ3, we compare the time cost for LM^{MR} as we increase the number of parallel tasks (see Fig. 8). We find that the performance improves as we increase the number of parallel tasks. By increasing the number of parallel MapReduce tasks from one to four, we gain a speed up in range of 1.4 to 3 times. By increasing the number of parallel MapReduce tasks from one to eight, we gain a speed up in range of 1.7 to 4.6 times. The performance improvement of LM^{MR} over LM is similar with that of CLIPPER^{MR} over CLIPPER. This is

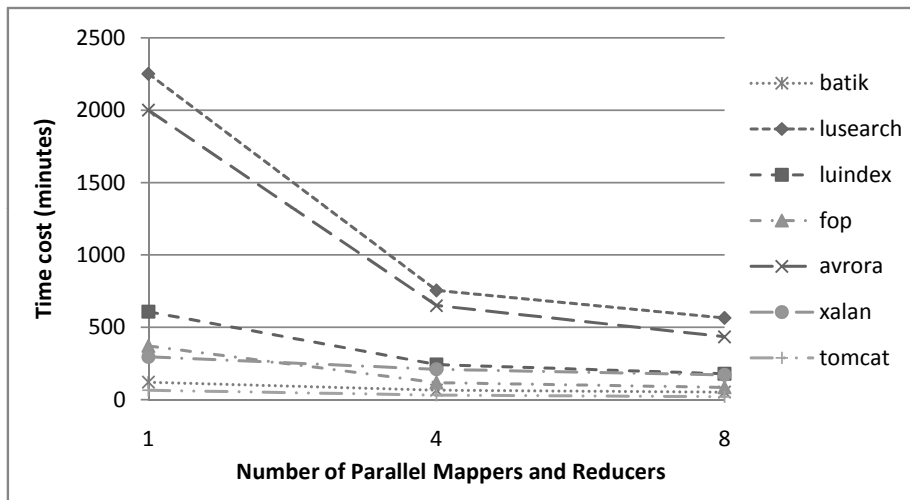


Figure 8: Performance Improvements for LM

reasonable as the first of the two steps of LM^{MR} is based on $CLIPPER^{MR}$ (see Section 3.2.4), and the second step for composing frequent charts into significant LSCs takes little time with respect to the time spent in the first step.

4.2.5. Mining Temporal Rules

To answer RQ1, we mine temporal rules with the original Perracotta, which was able to mine the temporal rules from all of the traces. Perracotta’s memory cost is quadratic to the number of unique events in the traces. In our study, the unique events are the methods that get invoked when program is run. The number of unique events is not so big, which is no more than 3000.

To answer RQ2, we compare the performance of the original Perracotta with that of $Perracotta^{MR}$ and present the results in Figure 9. We see that $Perracotta^{MR(8)}$ achieves speed up in range of 3.5 to 18.2. Note that $Perracotta^{MR(8)}$ may achieve more than 8 times speed up than Perracotta. This can be related to the fact that Perracotta is a memory-intensive algorithm (with a space complexity $O(n^2)$ and a time complexity $O(nL)$, where n is the number of unique events in the traces and L is the total length of all traces). Its sequential version needs to load all traces into memory sequentially, while the parallelized version may load many split smaller traces into memory simultaneously even when there is only one core available for map/reduce tasks. At the same time, the accuracy of $Perracotta^{MR}$ is 100% with respect to Perracotta: when we compare the output of $Perracotta^{MR}$ with that of Perracotta, there is no temporal rule that is missed.

To answer RQ3, we compare the time cost for $Perracotta^{MR}$ as we increase the number of parallel tasks (see Fig. 10). We find that the performance improves as we increase the number of parallel tasks. By increasing the number of parallel MapReduce tasks from one to four, we gain speed up in range of 1.6




Figure 9: Original versus Parallelized Perracotta




Figure 10: Performance Improvements for Perracotta

to 3.8 times. By increasing the number of parallel map reduce tasks from one to eight, we gain speed up in range of 4.1 to 7 times. Figure 10 also shows that the rate of speed up decreases as we increase the number of parallel tasks from 4 to 8. This is so as there are more resource contentions between the mappers and reducers as the number of parallel tasks are increased in our small four-machine cluster.

4.3. Threats to Validity and Current Limitations

In this work, we have considered five families of specification mining algorithms: those mining frequent patterns, value-based invariants, finite-state machines, sequence diagrams, and temporal rules. For each family, we have considered one algorithm. There are other specification mining algorithms that we do not consider in this study, e.g., those that analyze program code rather

than execution traces [34, 35, 36, 37, 26]. It is not clear if our approach can be easily extended to all other specification mining algorithms. In this study, we modify and adapt the algorithms to follow a divide-and-conquer strategy; it is not clear if all specification mining algorithms can be modified to follow this strategy. In the future, we would like to investigate even more algorithms and algorithm families and show how they can be modified to follow appropriate divide-and-conquer strategies to leverage the power of MapReduce.

We have evaluated our approach on seven programs from the DaCapo benchmark [7]. This benchmark has been used extensively in many past studies, e.g., [38, 39]. Still, these programs might not be representative of all open-source and industrial software systems. We plan to reduce this threat to validity further by investigating more programs in addition to those in the DaCapo benchmark in the future. Also, we have experimented with a cluster of four machines running eight cores. In the future, we plan to extend our experiment with even more machines and more cores. However, even with four machines, we have shown how the power of MapReduce could be tapped to scale various specification mining algorithms.

One limitation is imposed by the implementation of MapReduce that we use, i.e., Hadoop. One of the most important issues for a distributed platform is locality, as network bandwidth is the bottleneck when processing a large amount of data. To solve this problem, Hadoop attempts to replicate the data across the nodes and to always locate the nearest replica of the data. However, a substantial proportion of the time may still be spent on data transmission, typically for algorithms which involve heavy data transmission load. Our experiments use the default Hadoop file system replication factor 3 (i.e., each block of data is replicated to 3 machines) to minimize the transmission overheads during computation. The speedup factor of the parallelized versions of the specification mining algorithms may be affected if more machines are used or the replication factor is changed. In the future, we plan to perform a more comprehensive investigation of the effect of data transmission load, identify factors that may significantly affect the performance of distributed specification mining algorithms, and design improved algorithms that can reduce such overheads further.

5. Related Work

We now discuss closely related studies on specification mining, uses of MapReduce in software engineering, and parallelizing data mining algorithms in general. This section is by no means a complete list of related work.

5.1. Specification Mining and Its Applications

Mined specifications can help developers to understand legacy systems [8], to detect potential bugs [26]. They can also be used as input for model checkers and for the purpose of program verification [40] or be converted into test cases [41].

Some families of specification mining algorithms have been described in Section 2. Here we describe other related recent work. [42] mines three kinds

of temporal invariants from system logs and merges them into a state-based behavioral model. [43] proposes an approach that mines specifications from a variety of API data including information from API client programs, library source code, and comments. [44] uses three concepts, equivalence classes among LSCs, isomorphic embeddings, and delta-discriminative similarity measures, to mine a succinct set of LSCs and improve the readability of mining results. [45] presents a semi-automated approach that detects vacuously satisfiable scenarios leveraging a model checker and generates new scenarios to avoid the vacuity using machine learning (inductive logic programming). [4] presents a framework for mining message sequence graphs that can represent concurrent behaviors of a distributed system. [46] infers API specifications from documentations expressed in English. [47] implements a tool jMiner to mine parametric specifications by using a concept of trace slicing. The proposed approach first slices independent interactions from program traces. The independent interactions are then fed into a variant of the k-tails algorithm to produce a probabilistic finite-state machine. [48] builds on top of Daikon to infer contracts for Eiffel programs.

All five specification mining algorithms that we study in this chapter (CLIPPER, Daikon, k-tails, LM, and Perracotta) analyze program execution traces. Other approaches to specification mining use program code as input (e.g., [34, 35, 36, 37, 26]). A technique called *DySy*, mines invariants, similar to those generated by Daikon, by performing symbolic execution to reduce the number of test cases needed to mine the invariants and improve the quality of the mined invariants [49]. There are also algorithms that mine specifications in forms different from the five families described in Section 2, e.g, algebraic specifications [50] and separation logic invariants [51]. In the future, it would be interesting to examine how one can apply our general algorithm and MapReduce to the above diverse mining algorithms.

Some recent studies propose to better understand, extend, and compare existing specification mining algorithms. An approach called *InvariMint*, allows users to construct a model inference algorithm by using a declarative specification [52]. They show that their approach could help users understand, extend, and compare algorithms that mine specifications in the form of finite-state machines. Different from their work, we propose an approach to adapt existing algorithms to the MapReduce framework and make them more scalable and efficient.

A comprehensive survey of past studies on specification mining is available from a recent journal paper [53] and a book on specification mining [1].

5.2. MapReduce in Software Engineering

[12, 54] have presented experience reports on scaling tools for mining software repositories (MSR) using MapReduce. They investigate several case studies to analyze the potential of the MapReduce platform to scale up MSR tools, including (1) J-REX, which mines a CVS repository for calculating changes of software metrics over the history of a software project, (2) CC-Finder, which is a token-based clone detection tool designed to extract code clones from systems

developed in several programming languages, and (3) JACK, which is a log analyzer that uses data mining techniques to process system execution logs and automatically identify problems in load tests. Specification mining approaches are not covered in these studies. Recently, [11] propose a language and an infrastructure called Boa to ease the analysis of software repositories. Users can specify queries in a domain specific language and these queries can be processed by Boa’s processing engine which uses the MapReduce distributed computing model.

Different from these studies, we specifically focus on specification mining algorithms and investigate the potential of using MapReduce to make them more scalable.

5.3. Parallel Data Mining Algorithms

[55] uses MapReduce to propagate beliefs on a sparse billion-node graph. [56] uses MapReduce to parallelize an algorithm inferring document relevance for web search. [10] speeds up general clustering algorithms by using MapReduce. [9] recently proposes a parallelization of a frequent sequence mining algorithm that can run on MapReduce. Their approach relies on w -equivalency property that holds only for sequential patterns, and the semantics of the sequential patterns mined by their approach is different from that of iterative patterns mined by *CLIPPER* (which is the closest algorithm, considered in this chapter, to the frequent sequence mining algorithm). Although our approach employs MapReduce too, our data sources and the subject algorithms are specific to specification mining, which requires different parallelization strategies and techniques.

6. Conclusion and Future Work

In this chapter, we address the challenge of making specification mining algorithms scalable. We have presented a general algorithm design that helps to transform sequential specification mining algorithms into distributed ones based on the observation that many specification mining algorithms are data-intensive but computationally repetitive. In particular, we have shown how five different kinds of algorithms, *CLIPPER*, *Daikon*, *k-tails*, *LM*, and *Perracotta*, can be parallelized by following our general algorithm and leveraging a popular distributed computing framework MapReduce. We have evaluated the distributed versions of these algorithms with seven programs from the DaCapo benchmark and find that the distributed versions can significantly improve the scalability of the original algorithms for every trace dataset of sizes ranging from 41MB to 157GB. The distributed *Perracotta* running on four machines (using up to eight CPU cores in total) speeds up the original version by 3 to 18 times. The original *CLIPPER*, *Daikon*, *k-tails*, and *LM* are unable to handle the large traces, while our distributed versions can complete in hours and much performance improvements can be gained by utilizing more machines.

We consider the following for future work. First, our distributed algorithms are not necessarily optimal; we plan to investigate whether defining the map

and reduce functions differently and/or splitting input data differently would improve the scalability of these algorithms. For example, Daikon has more than one hundred invariant templates that are checked against traces when it looks for actual invariants. Checking each kind of template is independent from others, and thus could be parallelized as a map function as well. Second, our distributed algorithms are only evaluated with gigabyte traces in a four-machine cluster with many default settings; we would like to evaluate them with terabyte traces in a commercial cluster and see how performance improves when the number of processors increases and various cluster system settings are used. Third, we consider the application of our general algorithm and MapReduce to additional kinds of specification mining algorithms not covered in this chapter, e.g., algorithms leveraging other information besides traces (e.g., text, software repository). Fourth, some variants of the algorithms we have investigated may also deserve special attention, e.g., the variants of LSC mining triggers and effects [31], or the combination of scenario-based and value-based invariants [18].

- [1] D. Lo, S.-C. Khoo, J. Han, C. Liu (Eds.), *Mining Software Specifications: Methodologies and Applications*, CRC Press Data Mining and Knowledge Discovery Series, 2011.
- [2] D. Lo, S.-C. Khoo, C. Liu, Efficient mining of iterative patterns for software specification discovery, in: *KDD*, 2007, pp. 460–469.
- [3] M. D. Ernst, J. H. Perkins, P. J. Guo, S. McCamant, C. Pacheco, M. S. Tschantz, C. Xiao, The Daikon system for dynamic detection of likely invariants, *Sci. Comput. Program.* 69 (1-3) (2007) 35–45.
- [4] S. Kumar, S.-C. Khoo, A. Roychoudhury, D. Lo, Mining message sequence graphs, in: *ICSE*, 2011, pp. 91–100.
- [5] D. Lorenzoli, L. Mariani, M. Pezzè, Automatic generation of software behavioral models, in: *ICSE*, 2008, pp. 501–510.
- [6] D. Lo, S. Maoz, S.-C. Khoo, Mining modal scenario-based specifications from execution traces of reactive systems, in: *ASE*, 2007, pp. 465–468.
- [7] S. M. Blackburn, R. Garner, C. Hoffmann, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. L. Hosking, M. Jump, H. B. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanovic, T. VanDrunen, D. von Dincklage, B. Wiedermann, The DaCapo benchmarks: Java benchmarking development and analysis, in: *OOPSLA*, 2006, pp. 169–190.
- [8] J. Yang, D. Evans, D. Bhardwaj, T. Bhat, M. Das, Perracotta: mining temporal API rules from imperfect traces, in: *ICSE*, 2006, pp. 282–291.
- [9] I. Miliaraki, K. Berberich, R. Gemulla, S. Zoupanos, Mind the gap: large-scale frequent sequence mining, in: *SIGMOD Conference*, 2013, pp. 797–808.

- [10] A. Ene, S. Im, B. Moseley, Fast clustering using MapReduce, in: KDD, 2011, pp. 681–689.
- [11] R. Dyer, H. A. Nguyen, H. Rajan, T. N. Nguyen, Boa: A language and infrastructure for analyzing ultra-large-scale software repositories, in: Proceedings of the 2013 International Conference on Software Engineering, ICSE '13, 2013, pp. 422–431.
- [12] W. Shang, B. Adams, A. E. Hassan, An experience report on scaling tools for mining software repositories using MapReduce, in: ASE, 2010, pp. 275–284.
- [13] Message Passing Interface Forum, MPI: A message-passing interface standard (2012).
URL <http://www.mpi-forum.org/docs/docs.html>
- [14] J. Dean, S. Ghemawat, MapReduce: Simplified data processing on large clusters, in: OSDI, 2004, pp. 107–113.
- [15] Apache Software Foundation, Hadoop (Oct. 2013).
URL <http://hadoop.apache.org/>
- [16] J. Han, M. Kamber, Data Mining: Concepts and Techniques, Morgan Kaufman, 2006.
- [17] M. El-Ramly, E. Stroulia, P. G. Sorenson, From run-time behavior to usage scenarios: an interaction-pattern mining approach, in: KDD, 2002, pp. 315–324.
- [18] D. Lo, S. Maoz, Scenario-based and value-based specification mining: better together, in: ASE, 2010, pp. 387–396.
- [19] G. Ammons, R. Bodík, J. R. Larus, Mining specifications, in: POPL, 2002, pp. 4–16.
- [20] A. Biermann, J. Feldman, On the synthesis of finite-state machines from samples of their behavior, IEEE Transactions on Computers 21 (1972) 591–597.
- [21] L. C. Briand, Y. Labiche, J. Leduc, Toward the reverse engineering of UML sequence diagrams for distributed Java software, IEEE Trans. Software Eng. 32 (9) (2006) 642–663.
- [22] F. C. de Sousa, N. C. Mendonça, S. Uchitel, J. Kramer, Detecting implied scenarios from execution traces, in: WCRE, 2007, pp. 50–59.
- [23] W. Damm, D. Harel, LSCs: Breathing life into message sequence charts, Formal Methods in System Design (2001) 45–80.
- [24] D. Harel, S. Maoz, Assert and negate revisited: Modal semantics for UML sequence diagrams, Software and Systems Modeling 7 (2) (2008) 237–252.

- [25] V. B. Livshits, T. Zimmermann, Dynamine: finding common error patterns by mining software revision histories, in: ESEC/SIGSOFT FSE, 2005, pp. 296–305.
- [26] A. Wasylkowski, A. Zeller, Mining temporal specifications from object usage, *Autom. Softw. Eng.* 18 (3-4) (2011) 263–292.
- [27] D. Lo, S.-C. Khoo, C. Liu, Mining temporal rules for software maintenance, *Journal of Software Maintenance* 20 (4) (2008) 227–247.
- [28] Y.-F. Ho, S.-W. Chen, C.-Y. Chen, Y.-C. Hsu, P. Liu, A mapreduce programming framework using message passing, in: *International Computer Symposium (ICS)*, 2010, pp. 883–888.
- [29] Microsoft, Static driver verifier: DDI compliance rules.
URL [http://msdn.microsoft.com/en-us/library/ff552840\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ff552840(v=vs.85).aspx)
- [30] M. Pradel, T. R. Gross, Automatic generation of object usage specifications from large method traces, in: *ASE*, 2009, pp. 371–382.
- [31] D. Lo, S. Maoz, Mining scenario-based triggers and effects, in: *ASE*, 2008, pp. 109–118.
- [32] C. Chambers, A. Raniwala, F. Perry, S. Adams, R. R. Henry, R. Bradshaw, N. Weizenbaum, FlumeJava: Easy, efficient data-parallel pipelines, in: *PLDI*, 2010, pp. 363–375.
- [33] M. Gabel, Z. Su, Online inference and enforcement of temporal properties, in: *ICSE*, 2010, pp. 15–24.
- [34] Z. Li, Y. Zhou, PR-Miner: automatically extracting implicit programming rules and detecting violations in large software code, in: *ESEC/SIGSOFT FSE*, 2005, pp. 306–315.
- [35] T. T. Nguyen, H. A. Nguyen, N. H. Pham, J. M. Al-Kofahi, T. N. Nguyen, Graph-based mining of multiple object usage patterns, in: *ESEC/SIGSOFT FSE*, 2009, pp. 383–392.
- [36] S. Shoham, E. Yahav, S. Fink, M. Pistoia, Static specification mining using automata-based abstractions, in: *ISSTA*, 2007, pp. 174–184.
- [37] W. Weimer, G. C. Necula, Mining temporal specifications for error detection, in: *TACAS*, 2005, pp. 461–476.
- [38] M. D. Bond, K. E. Coons, K. S. McKinley, Pacer: proportional detection of data races, in: *PLDI*, 2010, pp. 255–268.
- [39] F. Chen, G. Rosu, Mop: an efficient and generic runtime verification framework, in: *OOPSLA*, 2007, pp. 569–588.

- [40] W. Li, A. Forin, S. A. Seshia, Scalable specification mining for verification and diagnosis, in: DAC, 2010, pp. 755–760.
- [41] V. Dallmeier, N. Knopp, C. Mallon, S. Hack, A. Zeller, Generating test cases for specification mining, in: ISSTA, 2010, pp. 85–96.
- [42] I. Beschastnikh, Y. Brun, S. Schneider, M. Sloan, M. D. Ernst, Leveraging existing instrumentation to automatically infer invariant-constrained models, in: SIGSOFT FSE, 2011, pp. 267–277.
- [43] Q. Wu, G.-T. Liang, Q.-X. Wang, H. Mei, Mining effective temporal specifications from heterogeneous API data, *J. Comput. Sci. Technol.* 26 (6) (2011) 1061–1075.
- [44] D. Lo, S. Maoz, Towards succinctness in mining scenario-based specifications, in: ICECCS, 2011, pp. 231–240.
- [45] D. Alrajeh, J. Kramer, A. Russo, S. Uchitel, Learning from vacuously satisfiable scenario-based specifications, in: FASE, 2012, pp. 377–393.
- [46] H. Zhong, L. Zhang, T. Xie, H. Mei, Inferring resource specifications from natural language API documentation, in: ASE, 2009, pp. 307–318.
- [47] C. Lee, F. Chen, G. Rosu, Mining parametric specifications, in: ICSE, 2011, pp. 591–600.
- [48] Y. Wei, C. A. Furia, N. Kazmin, B. Meyer, Inferring better contracts, in: ICSE, 2011, pp. 191–200.
- [49] C. Csallner, N. Tillmann, Y. Smaragdakis, DySy: Dynamic symbolic execution for invariant inference, in: ICSE, 2008, pp. 281–290.
- [50] J. Henkel, C. Reichenbach, A. Diwan, Developing and debugging algebraic specifications for Java classes, *ACM TOSEM* 17 (3) (2008) 14:1–14:37.
- [51] S. Magill, A. Nanevski, E. Clarke, P. Lee, Inferring invariants in separation logic for imperative list-processing programs, in: SPACE, 2006.
- [52] I. Beschastnikh, Y. Brun, J. Abrahamson, M. D. Ernst, A. Krishnamurthy, Unifying FSM-inference algorithms through declarative specification, in: ICSE, 2013, pp. 252–261.
- [53] M. P. Robillard, E. Bodden, D. Kawrykow, M. Mezini, T. Ratchford, Automated API property inference techniques, *IEEE Trans. Software Eng.* (2013) 613–637.
- [54] W. Shang, Z. M. Jiang, B. Adams, A. E. Hassan, MapReduce as a general framework to support research in mining software repositories (MSR), in: MSR, 2009, pp. 21–30.

- [55] U. Kang, D. H. Chau, C. Faloutsos, Mining large graphs: Algorithms, inference, and discoveries, in: ICDE, 2011, pp. 243–254.
- [56] C. Liu, F. Guo, C. Faloutsos, BBM: bayesian browsing model from petabyte-scale data, in: KDD, 2009, pp. 537–546.