

Automated Reasoning in the Simulation of Evolvable Systems

Djihed Afifi
University of Manchester
Manchester, U.K.
djihed@cs.man.ac.uk

David E. Rydeheard
University of Manchester
Manchester, U.K.
david@cs.man.ac.uk

Howard Barringer
University of Manchester
Manchester, U.K.
howard@cs.man.ac.uk

Abstract

We present a novel application of automated theorem proving for the simulation of computational systems. The computational systems we consider are *evolvable*, i.e. may reconfigure their structure and programs at run-time. In [1], a logical framework for describing such systems is introduced. The underlying logic of this framework allows us to build a simulation engine for executing system specifications. This engine makes intensive use of automated theorem proving – when running a simulation, almost all computational steps are those of a theorem prover. In this paper, we present this novel combination of a logical setting involving meta-level logics and large sets of formulae for system description, together with theorem proving requirements which involve often slowly changing specifications with the need for rapid assessment of deducibility and consistency. We will evaluate the suitability of several theorem provers for this application.

1 Introduction

We present a novel application of automated theorem proving for the simulation of computational systems. The computational systems we consider are *evolvable*, i.e. may reconfigure their structure and programs at run-time. In [1], a logical framework for describing such systems is introduced. The underlying logic of this framework allows us to build a simulation engine for executing system specifications. This engine makes intensive use of automated theorem proving – when running a simulation, almost all computational steps are those of a theorem prover. In this paper, we present this novel combination of a logical setting involving meta-level logics and large sets of formulae for system description, together with theorem proving requirements which involve often slowly changing specifications with the need for rapid assessment of deducibility and consistency. We will evaluate the suitability of several theorem provers for this application.

The architecture of evolvable systems that we employ allows the ‘localisation’ of monitoring and evolution. Components in a system may have ‘supervisors’ which are themselves components and which monitor their ‘supervisees’ and may evolve them if required. Such supervised components may be assembled hierarchically, with supervisors at each level of the hierarchy if necessary (compare Archware [9]).

In the logical framework, the state of a component is a set of ground atomic formulae which describe the current properties of the system. An action operates on the state in a style similar to STRIPS [5], changing the formula set by adding and deleting formulae. This revision-based method of description, which is common in planning and other AI applications, should be contrasted with the use pre and post-conditions [6]. This revision-based approach leads to a simple mechanical execution process, which we employ to build a simulator. This mechanisation makes intensive use of automated theorem proving (ATP) technology. Some issues relating to the appropriateness of the technology are:

- When a system undergoes an evolution, this may involve changing the current component theory as well as the component state. Thus as a system runs, theorem proving takes place in a highly dynamical setting of changing verification requirements in changing theories,

- Determining the applicability of actions: an executed action may have a set of logical formulae as preconditions. A valid application of an action requires that (a) the preconditions are derivable from the system state and the component theory, and (b) the resultant state is consistent.
- Handling meta-logics: the supervisor's state is at a *meta-level* to that of the supervisee allowing the supervisor to hold facts about the supervisee. the meta-level logic of the supervisor's state may hold facts about the supervisee. These facts need to be consistent in the supervisor-supervisee pairing at all times during execution.
- Verifiability for minimum models: The revision-based logic used in system specifications is based on a notion of 'minimum model' [1] and we require that deducibility is relative to this minimality requirement. This is related to the notions of the closed world assumption and circumscription in AI [8].

The simulations imposes a range of requirements on any theorem prover including the ability to:

- handle large sets of formulae for realistic systems,
- determine satisfiability of a set of formulae and deducibility of formulae from a given set of axioms.
- run unassisted, as opposed to guided, to make it possible to run large simulations that generate a high number of proof calls possible.
- construct models, not just satisfiability. Models are used to instantiate free variables in formulae.
- extract 'support sets' of formulae for proofs: States are often large sets of formulae, but those required to establish a property may be a small subset (often various different minimal 'support sets' may exist). Identifying support sets can aid proof caching (see below).

We have experimented with several theorem proving systems for this application:

- Paradox - from Chalmers University, a model finder for first-order logic [2].
- iProver - from The University of Manchester, an instantiation-based prover [7].
- Vampire - from The University of Manchester, a very fast resolution-based first-order theorem prover [10].

One point about the dynamics of theorem proving invocation needs explanation: during simulation, most individual actions cause a small change to the system, with the occasional evolution action that changes and reconfigures a system. On the other hand, actions generate multiple proof obligations that may duplicate prover dispatches. The execution of an action may trigger up to seven proof obligations for some action types. Furthermore, the overhead of discharging proof obligations is significant. The simulation of a simple system eventually spends a substantial amount of time communicating theories and proof results to and from the theorem prover. However, by examining simulations, we note that most changes do not affect state consistency, and in many cases proof obligations are duplicated or are trivially resolved.

In order to reduce the overhead of theorem proving various forms of proof caching have been employed. In general terms, proof caching is expensive. However, our particular application enables relatively simple approaches to speed up proof matching. The result is the elimination of a significant number of proof calls and dramatic speed-ups of simulations. In most cases more than 60% of proof obligations are eliminated with a corresponding increase in performance.




Figure 1: The scenario on the left shows the initial blocks world state, with a table having a capacity of 2 (i.e at most two blocks directly on the table) a tower of blocks. The scenario on the right shows a goal state that could be achieved by demolishing the tower and using a table of capacity 4.

This framework uses a non-standard type of inference relative to a minimum model determined by a set of ‘observable formulae’. This is a generalisation of the notion that absence of ground atoms indicates falsehood. The requirement to reason in ‘minimum models’ [1] is handled in simple cases by a completion process in which we close observable predicates. The general case of reasoning in minimum models combines this completion process with deducibility, but is not required for the system specifications we consider here.

In the next section, using an example, we will illustrate how theorem provers are used during the simulation of an evolvable system. We then describe a caching technique that eliminates a substantial number of proof dispatches. Finally, we discuss the results of experiments with several different types of theorem provers.

2 Logical Simulation of Evolvable Systems

We illustrate how theorem provers are used in the simulation with the following example. We start with the well-known blocks world [14]. This consists of number of blocks which may be on each other or on a table. Actions may move blocks around the world. A theory for a blocks world is defined in Figure 2. This is a simplified theory for illustration purposes. A complete axiomatisation of the blocks world is given in [3].

The blocks world theory *BlocksWorld* defines a finite set of blocks $\{A, B, C, D\}$ and a single table T . The ‘BWC’ axiom defines constraints on the predicate ‘on’. The constraint ‘TableSize’ is a parametric formula that initially defines tables with a capacity of 2, i.e at most two blocks may be directly on the table. The state of a blocks world component is recorded using ground atoms built from the binary predicate ‘on’. We distinguish between ‘observation’ predicates and ‘abstraction’ predicates. A state is described using only observation predicates. Only positive atoms of an observation predicate may be present in the state, those omitted are assumed to be false. The schema defines the abstraction predicate ‘free’ that will be used to deduce the availability of free space on the table. Abstraction predicates are defined using the constraints of the theory and may be deduced during simulation.

The component has the single action definition, ‘move’, that moves a block x from an object y to an object z . The action is conditional on the set ‘pre’. The action is performed on the state in a revision-based manner by deleting the atoms in the set ‘del’ and adding the atoms in the set ‘add’.

The component has an initial state where the blocks form a tower on top of the table as shown in Figure 1. The component is equipped with a program that demolishes the tower by moving all blocks on the table in turn. The program is a sequence of ‘move’ actions.

When this specification is executed, the first requirement is that it is *consistent*, meaning that its constraints and state together are consistent.

After checking the component’s consistency, the component’s program is run. The first action to be executed in the program is $move(D, C, T)$, which moves the the topmost block D in the initial state to the table T . The action move has the precondition set pre which checks that both the block and the

<i>BlocksWorld</i>									
<p>TYPES</p> $\text{Blocks} \stackrel{dfn}{=} \{A, B, C, D\}$ $\text{Tables} \stackrel{dfn}{=} \{T\}$ $\text{Objects} \stackrel{dfn}{=} \text{Blocks} \cup \text{Tables}$	<p>OBSERVATION PREDICATES</p> $\text{on} : \text{Blocks} \times \text{Objects}$ <p>ABSTRACTION PREDICATES</p> $\text{free} : \text{Objects}$								
<p>CONSTRAINTS</p> <p>$\text{BWC} \stackrel{dfn}{=} \forall b, b_1, b_2 : \text{Blocks}, o_1, o_2 : \text{Objects} \cdot$</p> $\neg \text{on}(b, b) \wedge$ $\text{on}(b, o_1) \wedge \text{on}(b, o_2) \Rightarrow (o_1 = o_2) \wedge$ $\text{on}(b_1, b_2) \Rightarrow (\exists o : \text{Objects} \cdot \text{on}(b_2, o))$ <p>$\text{TableSize}(T, 2) \stackrel{dfn}{=} (\exists b_1, b_2 : \text{Blocks} \cdot \text{on}(b_1, T) \wedge \text{on}(b_2, T) \wedge (b_1 \neq b_2)) \Leftrightarrow \neg \text{free}(T) \wedge$</p> $\forall b_1, b_2, b_3 : \text{Blocks} \cdot \text{on}(b_1, T) \wedge \text{on}(b_2, T) \wedge \text{on}(b_3, T) \Rightarrow ((b_1 = b_2) \vee (b_2 = b_3) \vee (b_1 = b_3))$ <p>$\text{BlockSize}(1) \stackrel{dfn}{=} \forall b : \text{Blocks} \cdot (\exists b_1 : \text{Blocks}, o : \text{Objects} \cdot \text{on}(b_1, b) \wedge \text{on}(b, o)) \Leftrightarrow \neg \text{free}(b) \wedge$</p> $\forall b_1, b_2 : \text{Blocks} \cdot \text{on}(b_1, b) \wedge \text{on}(b_2, b) \Rightarrow (b_1 = b_2)$									
<p>ACTIONS</p> <table border="1" style="width: 100%;"> <thead> <tr> <th>$\text{move}(x : \text{Blocks}, y, z : \text{Objects})$</th> <th></th> </tr> </thead> <tbody> <tr> <td>pre</td> <td>$\{\text{on}(x, z), \text{free}(x), \text{free}(y)\}$</td> </tr> <tr> <td>add</td> <td>$\{\text{on}(x, y)\}$</td> </tr> <tr> <td>del</td> <td>$\{\text{on}(x, z)\}$</td> </tr> </tbody> </table>		$\text{move}(x : \text{Blocks}, y, z : \text{Objects})$		pre	$\{\text{on}(x, z), \text{free}(x), \text{free}(y)\}$	add	$\{\text{on}(x, y)\}$	del	$\{\text{on}(x, z)\}$
$\text{move}(x : \text{Blocks}, y, z : \text{Objects})$									
pre	$\{\text{on}(x, z), \text{free}(x), \text{free}(y)\}$								
add	$\{\text{on}(x, y)\}$								
del	$\{\text{on}(x, z)\}$								
<p>INITIAL STATE</p> $\{\text{on}(D, C), \text{on}(C, B), \text{on}(B, A), \text{on}(A, T)\}$									
<p>PROGRAM</p> $\text{move}(D, C, T); \text{move}(C, B, T); \text{move}(B, A, T)$									

Figure 2: A Blocks World theory

receiving object are *free*. *Precondition checking* requires a theorem prover to establish deducibility of the preconditions from the state formulae and the component theory. If this is successful, the action revises the state by adding and deleting formulae as defined in the schema. The resultant state in turn needs to be checked for consistency with the component's theory to ensure that no action renders the component inconsistent.

In order to continue demolishing the tower an attempt is made at the second action in the program, $\text{move}(C, B, T)$. This will fail because the precondition cannot be met. The precondition $\text{free}(T)$ cannot be deduced as the current blocks world theory is restricted to just towers of blocks on the table as defined in *TableSize*. The only way to change this is to alter the theory. To achieve this, we introduce another system that monitors this system and has the ability to change the blocks world specification. We refer to this new system as a supervisor.

BlocksWorldSupervisor in Figure 4 presents such a supervisory system. It records properties of its *BlocksWorld* supervisee using meta-level predicates. The *holds* predicate is used to express properties of

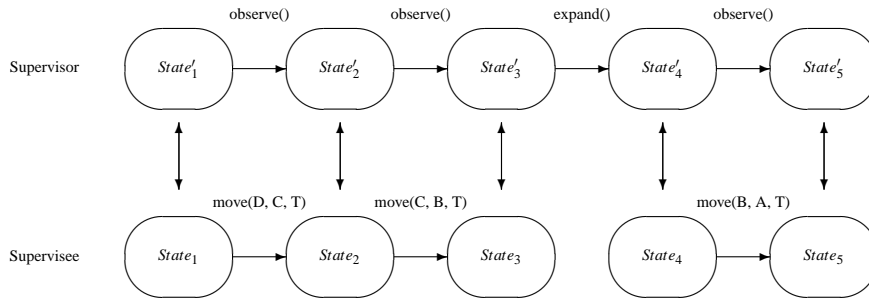


Figure 3: A paired execution trace

the supervisee. If $holds(\phi)$ occurs in the supervisor's state, then ϕ should be provable in the supervisee's state. The supervisor can use this predicate to query and monitor the supervisee. Similarly, the *constraint* predicate reflects the supervisee's constraints. The supervisor monitors the *BlocksWorld* using the action *observe* that tests formulae at the supervisee level. It also has the action *expand* which makes use of the evolution predicate *evolve*. The predicate *evolve* induces change at the supervisee level by revising the set of constraints that it has.

Using the *observe* action, the supervisor queries the state of the supervisee and detects when the table has an insufficient number of slots using the object level formula $free(T)$ as reflected at the supervisor level. This is a paired execution of the two example programs where the execution traces of the supervisor and the supervisee programs run in synchrony and are related by a meta-view relationship as depicted in Figure 3.

The supervisor monitors the supervisee and only intervenes when all all space on the table has been used, i.e when at the supervisee level $\neg free(T)$ holds. In this case, the *expand* action alters the *BlocksWorld* by replacing the *TableSize* constraint with one that allows for one more space. This is done using the *evolve* predicate, which introduces changes given by the supervisor to the supervisee. In this example, the change is the replacement of a constraint. The supervisor also has the ability to alter the state, redefine predicates or actions, or reconfigure the supervisee with a new component structure.

Meta-level conditions reflected by the predicate *holds* are checked by firing proof obligations using the supervisee's theory. The act of changing the supervisee's constraint using the *expand* action is called an *evolution step*. Theory consistency checks are necessary after performing the evolution for both components to disallow evolutions that lead to inconsistent theories.

Finally, component programs may have guarded instructions in which an instruction is executed only if its guard can be proved. We conclude this section by listing the cases where a theorem prover is invoked during the simulation of a component:

1. **Consistency** in the theory and state of each component, checked before and after each action. Models constructed in consistency checking may also be required.
2. **Precondition testing** by proving action preconditions.
3. **Guarded choice checking** by proving the guards.
4. **Meta-level checking** that the supervisee's reflection in the supervisor is correct, done before and after each evolution.
5. **Post-evolution consistency**, checking that evolutions don't produce inconsistent specifications.

In the above list, items 1 and 5 invoke theorem provers for satisfiability checking, while items 2, 3 and 4 invoke derivability checks.

<i>BlocksWorldSupervisor</i> META TO <i>BlocksWorld</i>																	
<p>TYPES <i>ConfigName</i></p> <p>FUNCTIONS $s : \text{ConfigName} \rightarrow \text{ConfigName}$</p>	<p>OBSERVATION PREDICATES <i>current</i> : <i>ConfigName</i> <i>holds</i> : FORMULA \times <i>ConfigName</i> <i>constraint</i> : CONSTRAINTNAME <i>evolve</i> : ATOMS \times ATOMS \times CONSTRAINTNAMES \times CONSTRAINTNAMES \times <i>ConfigName</i></p>																
<p>CONSTRAINTS ...</p> <p>ACTIONS</p> <table border="1" style="width: 100%;"> <tr> <td style="padding: 2px;"><i>observe</i>($P : \text{FORMULAE}$)</td> <td style="border-left: 1px solid black;"></td> </tr> <tr> <td style="padding: 2px;">pre</td> <td style="border-left: 1px solid black;">$\{ \text{current}(c) \}$</td> </tr> <tr> <td style="padding: 2px;">add</td> <td style="border-left: 1px solid black;">$\{ \text{holds}(p, s(c)) \mid p \in P \} \cup \{ \text{current}(s(c)) \}$</td> </tr> <tr> <td style="padding: 2px;">del</td> <td style="border-left: 1px solid black;">$\{ \text{current}(c) \}$</td> </tr> <tr> <td style="padding: 2px;"><i>expand</i>($n : \text{Int}$)</td> <td style="border-left: 1px solid black;"></td> </tr> <tr> <td style="padding: 2px;">pre</td> <td style="border-left: 1px solid black;">$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)), m < n \}$</td> </tr> <tr> <td style="padding: 2px;">add</td> <td style="border-left: 1px solid black;">$\{ \text{current}(s(c)), \text{holds}(\text{free}(T), s(c)), \text{evolve}(\{ \}, \{ \}, \{ \text{TableSize}(T, n) \}, \{ \text{TableSize}(T, m) \}), \{ \}, \{ \}, s(c), \text{constraint}(\text{TableSize}(T, n)) \}$</td> </tr> <tr> <td style="padding: 2px;">del</td> <td style="border-left: 1px solid black;">$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)) \}$</td> </tr> </table>	<i>observe</i> ($P : \text{FORMULAE}$)		pre	$\{ \text{current}(c) \}$	add	$\{ \text{holds}(p, s(c)) \mid p \in P \} \cup \{ \text{current}(s(c)) \}$	del	$\{ \text{current}(c) \}$	<i>expand</i> ($n : \text{Int}$)		pre	$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)), m < n \}$	add	$\{ \text{current}(s(c)), \text{holds}(\text{free}(T), s(c)), \text{evolve}(\{ \}, \{ \}, \{ \text{TableSize}(T, n) \}, \{ \text{TableSize}(T, m) \}), \{ \}, \{ \}, s(c), \text{constraint}(\text{TableSize}(T, n)) \}$	del	$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)) \}$	
<i>observe</i> ($P : \text{FORMULAE}$)																	
pre	$\{ \text{current}(c) \}$																
add	$\{ \text{holds}(p, s(c)) \mid p \in P \} \cup \{ \text{current}(s(c)) \}$																
del	$\{ \text{current}(c) \}$																
<i>expand</i> ($n : \text{Int}$)																	
pre	$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)), m < n \}$																
add	$\{ \text{current}(s(c)), \text{holds}(\text{free}(T), s(c)), \text{evolve}(\{ \}, \{ \}, \{ \text{TableSize}(T, n) \}, \{ \text{TableSize}(T, m) \}), \{ \}, \{ \}, s(c), \text{constraint}(\text{TableSize}(T, n)) \}$																
del	$\{ \text{current}(c), \text{constraint}(\text{TableSize}(T, m)) \}$																
<p>INITIAL STATE $\{ \text{current}(c_0) \}$</p>																	

Figure 4: The Blocks World Supervisor schema

2.1 System Simulation

In the example, the pairing of *BlocksWorldSupervisor* and *BlocksWorld* specifications will, when executed, ensure the complete demolition of the tower. The supervisor intervenes repeatedly to evolve the supervisee and expand the table whenever it cannot accommodate any more blocks. To achieve this, the supervisor is equipped with the following program:

```
([
  |observe(free(T))
  |(observe(¬free(T));expand())
])*;
is_demolished();
```

This program is an iteration of a non-deterministic choice instruction that queries whether the table is free after every supervisee move action. The *expand* action is only performed when the table is observed to be not free. The *is_demolished* action performs a logical check to determine whether the demolition objective was successfully reached.

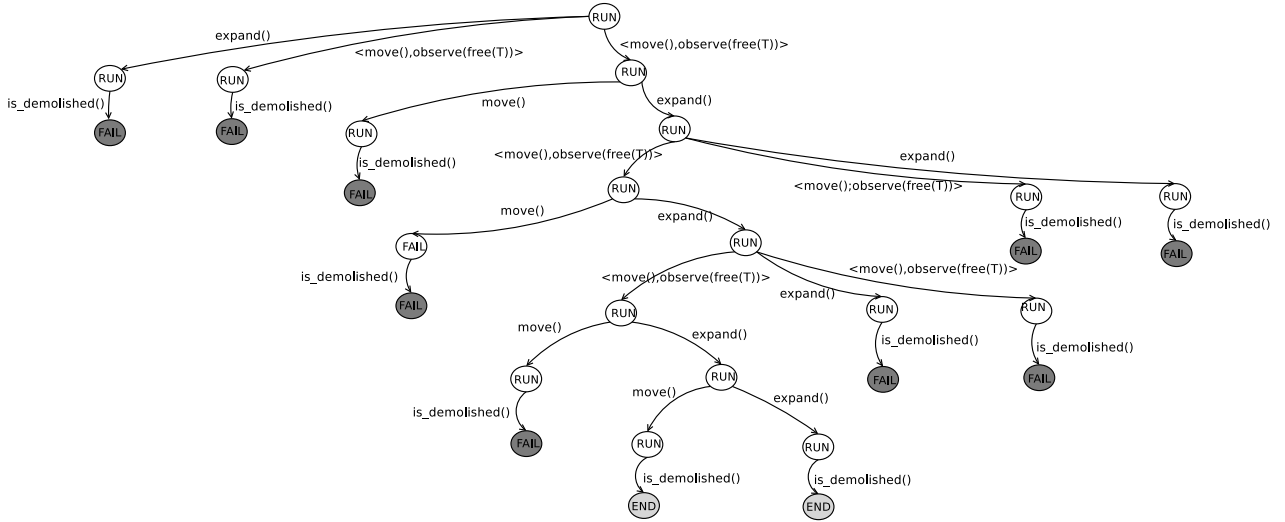


Figure 5: A complete run of the blocks world

A complete run of this paired simulation will result in a configuration trace tree as shown in figure 5. Each node in this tree represents a state of the system configuration. This simulation is exhaustive as it looks for all possible runs of the pair of the components (this search can either be depth-first or breadth-first). The dark nodes indicate actions that have failed: this occurs when a supervisee’s *move* action is not possible, or when an evolution is necessary. The grey nodes represent successful runs where the run successfully demolished the tower. Following the nodes in the trace tree, these successful runs occur when each block move is immediately preceded by an evolution to expand the table size.

3 Caching and Eliminating Proof Obligations

Simulating specifications can generate a substantial number of proof obligations. A basic non-supervisory action generates two proof obligations for every sub-component that it modifies, but an evolution action generates seven proof obligations for each pair of components that it affects. The simulation of the previous relatively simple blocks world generates 35 deducibility checks and 29 satisfiability checks. Firing an external theorem prover and communicating theories and results uses a substantial amount of time relative to the simulation execution time. Given that in any realistic simulation very large states will result and that the performance of a theorem prover is typically non-linear with the size of the state, it is important to reduce the number of dispatched proof obligations. Several techniques can be used to optimise execution and discharge some of the proof obligations without the need to call an external theorem prover.

By examining simulations, we note that the most commonly executed actions are of a basic type, which are occasionally interrupted by the rarer large system reconfigurations that alter component theories or replace components. Most actions do not affect the axioms of a component and therefore do not affect its internal consistency. A substantial number of consistency checks can therefore be eliminated by storing the consistency results of previous prover invocations. A component is given a ‘consistent’ flag that is set once and is reset when actions changes are deemed to affect this consistency.

The state of a component contains only positive atoms of observation predicates. Therefore, in some cases, when preconditions are themselves ground observable atoms, derivability reduces to testing

<i>Cache</i>	$\stackrel{dfn}{=}$	$ComponentCacheKey \rightarrow ComponentCache$
<i>ComponentCacheKey</i>	$\stackrel{dfn}{=}$	$SchemaID \times ConstraintNames$
<i>ConstraintNames</i>	$\stackrel{dfn}{=}$	$ConstraintName^*$
<i>ComponentCache</i>	$\stackrel{dfn}{=}$	$State \rightarrow ProvenFormulae \times UnprovenFormulaeSets$
<i>State</i>	$\stackrel{dfn}{=}$	$Atom^*$
<i>ProvenFormulae</i>	$\stackrel{dfn}{=}$	$Formula^*$
<i>UnprovenFormulaeSet</i>	$\stackrel{dfn}{=}$	$FormulaeSet^*$
<i>FormulaeSet</i>	$\stackrel{dfn}{=}$	$Formula^*$

Figure 6: Cache Structure

membership using symbolic equality. The minimum model interpretation means that the absence of an atom of an observation predicates indicates its falsehood. This allows us to reduce preconditions that rely heavily on observable predicates. It is often the case that the precondition set is reduced to either the empty set (i.e preconditions are met) or to *false*, eliminating the need for any external theorem prover.

3.1 Caching of Prover Results

A cache structure suitable for this application is depicted in Figure 6. Properties of our framework make this cache structure relatively efficient.

The cache stores the prover invocation results for each component separately. To perform lookup, a key for the cache consists of the name of the component together with the names of its constraints. Only the signatures of the constraints are stored, e.g. $BlockSize(2, T)$. This eliminates the need to store whole formulae and simplifies component cache lookup. An example of a cache key is

$$(BlocksWorld, \{TableSize(2, T), BlockSize(1)\})$$

Each component is then associated with all states that it may have had in the past. For each state that a component may have been in, the list of proved formulae is stored. Formulae lookups are done syntactically, so this is efficient for looking up previously proved ground atoms. The cache also stores the sets of formulae that were previously disproved given a state.

Looking up formulae in this cache is undertaken as follows. Given a proof obligation for the formulae set S , once a cache key has been matched together with a state, S is reduced to the set S' by eliminating formulae that were previously proved. If S' is the empty set then this is considered a cache hit as proved formulae and a prover invocation is not necessary. if S' is not empty, it is compared with every set of the previously unproven formulae. If S' contains any of these sets then this is also a cache hit as unproven formulae. If S' is not matched by any unproven formulae set, an external prover invocation is necessary to determine the deducibility of S' . If the prover returns a proved result then S' is merged with the proved formulae in the cache. if S' is not proved it is added to the set of previously unproven formulae sets.

3.2 Performance

The caching technique eliminates a substantial number of proof obligations. In the blocks world example, 20 out of 35 deduction requests (57%) and 24 out of 29 satisfiability requests (82%) are eliminated. The running time is reduced from 16 seconds to 5 seconds on an AMD Athlon 2000+ processor with 1GB of RAM. Figure 7 shows the number of total proof dispatches and actual external prover invocations for specifications that generate increasing program traces. The trace size is an indication of the number of

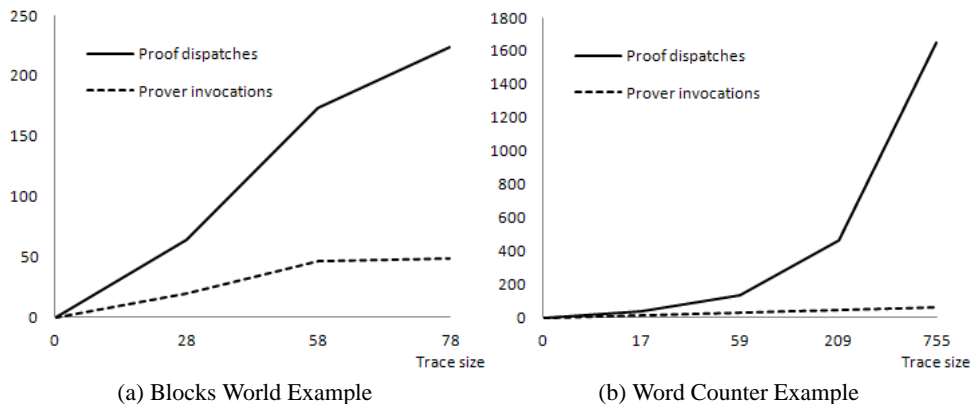


Figure 7: Prover invocation elimination using the caching technique

actions being performed. The figure shows that the longer the simulation the more beneficial caching becomes, with some simulations eliminating over 90% of all proof obligations.

4 Comparing Theorem Provers

The simulator converts theories to classical untyped first order logic in the TPTP3 [13] format so CASC (CADE ATP System Competition) [12] can be used. Three provers were used in this study:

- **Paradox** [2]: a finite model generator that flattens first order logic formulae into proposition clauses, then uses the MiniSat [4] SAT solver to solve the resulting problem.
- **Vampire 10** [10]: a fast resolution-based theorem prover.
- **iProver** [7]: an instantiation based prover that combines first order reasoning with a SAT solver (using MiniSat [4]).

The table in Figure 8 shows the running time in seconds of the simulation using these provers. The measure of the complexity of a simulation is the trace size, which grows approximately linearly with the number of proof obligations being made.

For the blocks world example, Paradox was able to undertake both the satisfiability checks and the proofs. It is fast at finding models for axiom sets and for checking counter satisfiability [11]. Overall, it was the best performer of the three theorem provers for this application. Its ability to determine the deducibility or otherwise the counter satisfiability of formulae makes it the most suitable choice for this application.

Although Vampire’s resolution is fast, it is not appropriate to establish the counter satisfiability of non-theorems, nor could it establish the satisfiability of theories in the above example. This necessitates the use of other theorem provers for these purposes when using Vampire.

iProver is unique because it can do both resolution reasoning and SAT checking, but they are done successively with manual options to turn off either features. Its resolution reasoning is fast at establishing the non-satisfiability of sets of formulae and for deriving theorems, while its SAT solving mode is fast

Trace size	Paradox	iProver	Vampire
28	5	12	15
58	8	18	22
78	11	24	32

Figure 8: Simulation time (in seconds) using different provers. Note that satisfiability checking was done with Paradox when using Vampire.

at establishing the satisfiability of sets of axioms and the counter-satisfiability of non-theorems. iProver can spend time unnecessarily using one of its modes for an input that is best suited for the other mode.

5 Conclusion

This paper gave an overview of the practical aspects of using theorem provers for the simulation of the evolvable systems framework presented in [1]. In this simulation, theorem provers are used to deduce formulae from an axiom set and for establishing the satisfiability of sets of formulae. The simulation generates a large number of proof obligations. However, a caching technique was used to eliminate a substantial number of prover invocations and speed up simulation.

Three theorem provers were used in this study: Paradox [2], iProver [7] and Vampire [10]. Paradox's model finding was best suited for establishing satisfiability of axiom sets as well as the counter satisfiability of non-theorems. Although vampire's resolution is fast, it could not establish the counter-satisfiability of non-theorems. iProver's manually changeable two modes, resolution and SAT solving, could perform both tasks. In the future, running multiple theorem provers in parallel could be performed to exploit the strength of each prover.

References

- [1] H. Barringer, D. Gabbay, and D. E. Rydeheard. Modelling evolvable component systems. Part I. A logical framework. *Logic Jnl IGPL*, page jzp026, 2009.
- [2] K. Claessen and N. Sörensson. New techniques that improve MACE-style model finding. *Proc. of Workshop on Model Computation (MODEL)*, 2003.
- [3] S.A. Cook and Y. Liu. A complete axiomatization for blocks world. *Journal of Logic and Computation*, 13(4):581, 2003.
- [4] Niklas Eén and Niklas Sörensson. An extensible sat-solver. Santa Margherita Ligure, Italy, May 2003.
- [5] R.E. Fikes and N.J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4):189–208, 1971.
- [6] C. B. Jones. *Systematic Software Development using VDM*. International Series in Computer Science. Prentice-Hall, 1990.
- [7] K. Korovin. iProver - an instantiation-based theorem prover for first-order logic (system description). *IJCAR*, volume 5195 of *Lecture Notes in Computer Science*, pages 292–298. Springer, 2008.
- [8] J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial intelligence*, 13(1-2):27–39, 1980.
- [9] F. Oquendo, B. Warboys, R. Morrison, R. Dindeleux, F. Gallo, H. Garavel, and C. Occhipinti. ArchWare: Architecting Evolvable Software. *Proceedings of the 1st European Workshop on Software Architecture (EWSA'04), European Projects in Software Architecture - Invited Paper*, LNCS, pages 257–271, St Andrews, UK, May 2004.
- [10] A. Riazanov and A. Voronkov. The design and implementation of VAMPIRE. *AI Communications*, 15(2-3):91–110, 2002.
- [11] G. Sutcliffe. The 4th IJCAR Automated Theorem Proving Competition. *AI Communications*, 22(1):59–72, 2009.
- [12] G. Sutcliffe and C. Suttner. The State of CASC. *AI Communications*, 19(1):35–48, 2006.
- [13] G. Sutcliffe and C.B. Suttner. The TPTP Problem Library: CNF Release v1.2.1. *Journal of Automated Reasoning*, 21(2):177–203, 1998.
- [14] T Winograd. *Understanding Natural Language*. Academic Press, New York, 1972.