MONA Implementation Secrets

Nils Klarlund[1], Anders Møller[2], and Michael I. Schwartzbach[2]

[1] AT&T Labs–Research, klarlund@research.att.com
[2] BRICS, University of Aarhus, {amoeller,mis}@brics.dk

Abstract. The Mona tool provides an implementation of the decision procedures for the logics WS1S and WS2S. It has been used for numerous applications, and it is remarkably efficient in practice, even though it faces a theoretically non-elementary worst-case complexity. The implementation has matured over a period of six years. Compared to the first naive version, the present tool is faster by several orders of magnitude. This speedup is obtained from many different contributions working on all levels of the compilation and execution of formulas. We present a selection of implementation "secrets" that have been discovered and tested over the years, including formula reductions, DAGification, guided tree automata, three-valued logic, eager minimization, BDD-based automata representations, and cache-conscious data structures. We describe these techniques and quantify their respective effects by experimenting with separate versions of the Mona tool that in turn omit each of them.

1 Introduction

Mona [14, 20, 22] is an implementation of the decision procedures for the logics WS1S and WS2S [28]. They have long been known to be decidable [7, 8, 13], but with a non-elementary lower bound [21]. For many years it was assumed that this discouraging complexity precluded any useful implementations.

However, Mona has been developed at BRICS since 1994, when our initial attempt at automatic pointer analysis through automata calculations took four hours to complete. Today Mona has matured into an efficient and popular tool on which the same analysis is performed in a couple of seconds. Through those years, many different approaches have been tried out, and a good number of implementation "secrets" have been discovered. This paper describes the most important tricks we have learned, and it tries to quantify their relative merits on a number of benchmark formulas.

Of course, the resulting tool still has a non-elementary worst-case complexity. Perhaps surprisingly, this complexity also contributes to successful applications, since it is provably linked to the succinctness of the logics. If we want to describe a particular regular set, then a WS1S formula may be non-elementarily more succinct that a regular expression or a transition table.

The niche for Mona applications contains those structures that are too large and complicated to describe by other means, yet not so large as to require infeasible computations. Happily, many interesting projects fit into this niche, including hardware verification [2, 3], pointer analysis [16, 12], controller synthesis [25, 15], natural languages [23], parsing tools [10], Presburger arithmetic [26], and verification of concurrent systems [17, 1, 24, 27].

2 MONA, WS1S, and WS2S

The first versions of Mona were based on a logic about finite strings, the monadic second-order logic M2L(Str). In this notation, first-order variables are interpreted over the positions in a finite string. Thus, for a given interpretation, there is a maximum value that a first-order variable may take on. A second-order variable denotes a subset of positions. A formula is valid if it holds for any finite string. The decision procedure for this logic is slightly easier to implement than that of the current Mona, which is based on WS1S. This logic is simpler to explain: a first-order variable denotes a natural number, and a second-order variable denotes a finite set of numbers. Both logics allow the comparison of variables in the expected ways depending on their order: $<$, $\subseteq$, $=$, $\in$, etc. Also, a function symbol $+1$ is allowed on first-order terms. It denotes the successor (where in the case of M2L(Str), the successor of the last position is defined as the first position).

MONA additionally supports the logic WS2S with two successors. Also, there is explicit syntax for Presburger constants. Finally, it implements the variation WSRT which allows values of recursive data types rather than simply binary trees. The MONA manual [20] describes the syntax and semantics of the MONA language and the features of the tool.

The automaton for a formula is constructed recursively from automata representing subformulas. In the cases of M2L(Str) and WS1S, each automaton describes a language of strings over the alphabet $\{0,1\}^k$, where k is the number of free variables in the subformula. Each string represents an interpretation, that is, an assignment of values to variables that are free in the subformula; the language is the set of strings that define satisfying interpretations. For WS2S this is generalized to tree automata.

3 Benchmark Formulas

The experiments presented in the following section are based on twelve benchmark formulas, here shown with their sizes, the logics they are using, and their time and space consumptions when processed by MONA 1.4 (on a 296MHz UltraSPARC with 1GB RAM):

Benchmark	Name	Size	Logic	Time	Space
A	`dflipflop.mona`	2 KB	WS1S (M2L-Str)	0.4 sec	3 MB
B	`euclid.mona`	6 KB	WS1S (Presburger)	33.1 sec	217 MB
C	`fischer_mutex.mona`	43 KB	WS1S	15.1 sec	13 MB
D	`html3_grammar.mona`	39 KB	WS2S (WSRT)	137.1 sec	208 MB
E	`lift_controller.mona`	36 KB	WS1S	8.0 sec	15 MB
F	`mcnc91_bbsse.mona`	9 KB	WS1S	13.2 sec	17 MB
G	`reverse_linear.mona`	11 KB	WS1S (M2L-Str)	3.2 sec	4 MB
H	`search_tree.mona`	19 KB	WS2S (WSRT)	30.4 sec	5 MB
I	`sliding_window.mona`	64 KB	WS1S	40.3 sec	59 MB
J	`szymanski_acc.mona`	144 KB	WS1S	20.6 sec	9 MB
K	`von_neumann_adder.mona`	5 KB	WS1S	139.9 sec	116 MB
L	`xbar_theory.mona`	14 KB	WS2S	136.4 sec	518 MB

The benchmarks have been picked from a large variety of MONA applications ranging from hardware verification to encoding of natural languages.

`dflipflop.mona` – a verification of a D-type flip-flop circuit [3]. Provided by Abdel Ayari.

`euclid.mona` – an encoding in Presburger arithmetic of six steps of reachability on a machine that implements Euclid's GCD algorithm [26]. Provided by Tom Shiple.

`fischer_mutex.mona` and `lift_controller.mona` – duration calculus encodings of Fischer's mutual exclusion algorithm and a mine pump controller, translated to MONA code [24]. Provided by Paritosh Pandya.

`html3_grammar.mona` – a tree-logic encoding of the HTML 3.0 grammar annotated with 10 parse-tree formulas [10]. Provided by Niels Damgaard.

`reverse_linear.mona` – verifies correctness of a C program reversing a pointer-linked list [16].

`search_tree.mona` – verifies correctness of a C program deleting a node from a search tree [12].

`sliding_window.mona` – verifies correctness of a sliding window network protocol [27]. Provided by Mark Smith.

`szymanski_acc.mona` – validation of the parameterized Szymanski problem using an accelerated iterative analysis [5]. Provided by Mamoun Filali-Amine.

`von_neumann_adder.mona` and `mcnc91_bbsse.mona` – verification of sequential hardware circuits; the first verifies that an 8-bit von Neumann adder is equivalent to a standard carry-chain adder, the second is a benchmark from MCNC91 [29]. Provided by Sebastian Mödersheim.

`xbar_theory.mona` – encodes a part of a theory of natural languages in the Chomsky tradition. It was used to verify the theory and led to the discovery of mistakes in the original formalization [23]. Provided by Frank Morawietz.

We will use these benchmarks to illustrate the effects of the various implementation "secrets" by comparing the efficiency of MONA shown in the table above with that obtained by handicapping the MONA implementation by not using the techniques.

4 Implementation Secrets

The MONA implementation has been developed and tuned over a period of six years. Many large and small ideas have contributed to a combined speedup of several orders of magnitude. Improvements have taken place at all levels, which we illustrate with the following seven examples from different phases of the compilation and execution of formulas.

To enable comparisons, we summarize the effect of each implementation "secret" by a single dimensionless number for each benchmark formula. Usually, this is simply the speedup factor, but in some cases where the numerator is not available, we argue for a more synthetic measure. If a benchmark cannot run on our machine, it is assigned time ∞.

4.1 Eager minimization

When MONA inductively translates formulas to automata, a Myhill-Nerode minimization is performed after every product and projection operation. Naturally, it is preferable to operate with as small automata as possible, but our strategy may seem excessive since minimization often exceeds 50% of the total running time. This suspicion is strengthened by the fact that MONA automata by construction contain only reachable states; thus, minimization only collapses redundant states.

Three alternative strategies to the eager one currently used by MONA would be to perform only the very final minimization, only the ones occurring after projection operations, or only the ones occurring after product operations. Many other heuristics could of course also be considered. The following table results from such an investigation:

Benchmark	Time				Effect
	Only final	After project	After product	Always	
A	∞	∞	0.6 sec	0.4 sec	1.5
B	∞	∞	∞	33.1 sec	∞
C	∞	∞	32.3 sec	15.1 sec	2.1
D	∞	∞	290.6 sec	137.1 sec	2.1
E	∞	∞	19.4 sec	8.0 sec	2.4
F	∞	∞	36.7 sec	13.2 sec	2.8
G	∞	∞	5.8 sec	3.2 sec	1.8
H	∞	∞	59.6 sec	30.4 sec	2.0
I	∞	∞	74.4 sec	40.3 sec	1.8
J	∞	∞	36.3 sec	20.6 sec	1.8
K	∞	∞	142.3 sec	139.9 sec	1.0
L	∞	∞	∞	136.4 sec	∞

"Only final" is the running time when minimization is only performed as the final step of the translation; "After project" is the running time when minimization is also performed after every projection operation; "After product" is the running time when minimization is instead performed after every product operation; "Always" is the time when minimization is performed eagerly; and "Effect" is the "After product" time compared to the "Always" time (since the other two strategies are clearly hopeless). Eager minimization is seen to be always beneficial and in some cases essential for the benchmark formulas.

4.2 Guided tree automata

Tree automata are inherently more computationally expensive because of their three-dimensional transition tables. We have used a technique of factorization of state spaces to split big tree automata into smaller ones. The basic idea, which may result in exponential savings, is explained in [4]. To exploit this feature, the MONA programmer must manually specify a *guide*, which is a top-down tree automaton that assigns state spaces to the nodes of a tree. However, when using the WSRT logic, a canonical guide is automatically generated. For our two WSRT benchmarks, we measure the effect of this canonical guide:

Benchmark	Without guide	With guide	Effect
D	584.0 sec	137.1 sec	4.3
H	∞	30.4 sec	∞

"Without guide" shows the running time without any guide, while "With guide" shows the running time with the canonical WSRT guide; "Effect" shows the "Without guide" time compared to the "With guide" time. We have only a small sample space here, but clearly guides are very useful. This is hardly surprising, since they may yield an asymptotic improvement in running time.

4.3 Cache-conscious data structures

The data structure used to represent the BDDs for transition functions has been carefully tuned to minimize the number of cache misses that occur. This effort is motivated in earlier work [18], where it is determined that the number of cache misses during unary and binary BDD apply steps totally dominates the running time.

In fact, we argued elsewhere [18] that if A1 is the number of unary apply steps and A2 is the number of binary apply steps, then there exists constant m, c_1, and c_2 such that the total running time is approximately $m(c_1 \cdot \mathrm{A1} + c_2 \cdot \mathrm{A2})$. Here, m is the machine dependent delay incurred by an L2 cache miss, and c_1 and c_2 are the average number of cache misses for unary and binary apply steps. This estimate is based the assumption that time incurred for manipulating auxiliary data structures, such as those used for describing subsets in the determinization construction, is insignificant. For the machine we have used for experiments, it is by a small C utility determined that $m = 0.43\mu s$. In our BDD implementation, explained in [18], we have estimated from algorithmic considerations that $c_1 = 1.7$ and $c_2 = 3$ (the binary apply may entail the use of unary apply steps for doubling tables that were too small—these steps are not counted towards the time for binary apply steps, and that is why we can use the figure $c_2 = 3$); we also estimated that for an earlier conventional implementation, the numbers were $c_1 = 6.7$ and $c_2 = 7.3$. The main reason for this difference is that our specialized package stores nodes directly under their hash address to minimize cache misses; traditional BDD packages store BDD nodes individually with the hash table containing pointers to them—roughly doubling the time it takes to process a node. We no longer support the conventional BDD implementation, so to measure the effect of cache-consciousness, we must use the above formula to estimate the running times that would have been obtained today.

In the following experiment, we have instrumented Mona to obtain the exact numbers of apply steps:

Benchmark	Apply1	Apply2	Misses	Auto	Predicted	Conventional	Effect
A	183,949	28,253	397,472	0.2 sec	0.2 sec	0.6 sec	3.0
B	21,908,722	3,700,856	48,347,395	32.8 sec	20.8 sec	74.7 sec	3.6
C	24,585,292	1,428,381	46,080,139	14.2 sec	19.8 sec	75.2 sec	3.8
E	9,847007	822,796	19,208,299	7.7 sec	8.2 sec	30.9 sec	3.8
F	13,406,047	5,717,453	39,942,638	12.8 sec	17.2 sec	56.6 sec	3.3
G	233,566	54,814	561,504	0.5 sec	0.3 sec	0.8 sec	2.7
I	36,629,195	11,153,733	95,730,831	37.0 sec	41.2 sec	140.5 sec	3.4
J	10,497,759	2,257,791	24,619,563	11.6 sec	10.6 sec	37.3 sec	3.5
K	129,126,447	10,485,623	250,971,828	137.4 sec	107.9 sec	404.7 sec	3.8

"Apply1" is the number of unary apply steps; "Apply2" is the number of binary apply steps; "Misses" is the number of cache misses predicted by the formula above; "Auto" is the part of the actual running time involved in automata constructions; "Predicted" is the running time predicted from the cache misses alone; "Conventional" is the predicted running time for a conventional BDD implementation that was not cache-conscious; and "Effect" is "Conventional" compared to "Predicted". In most cases, the actual running time is close to the predicted one (within 25%). Note that there are instances where the actual time is about 50% larger than the estimated time: benchmark B involves a lengthy subset construction on an automaton with small BDDs—thus it violates the assumption that the time handling accessory data structures is insignificant; similarly, benchmark G also consists of many automata with few BDD nodes prone to violating the assumption.

In an independent comparison [26] it was noted that MONA was consistently twice as fast as a specially designed automaton package based on a BDD package considered efficient. In [18], the comparison to a traditional BDD package yielded a factor 5 speedup.

4.4 BDD-based automata representation

Its reasonable to ask: "What would happen if we had simply represented the transition tables in a standard fashion, that is, a row for each state and a column for each letter?". Under this point of view, it makes sense to define a letter for each bit-pattern assignment to the free variables of a subformula (as opposed to the larger set of all variables bound by an outer quantifier). We have instrumented MONA to measure the sum of the number of entries of all such automata transition tables constructed during a run of a version of MONA without BDDs:

Benchmark	Misses	Table entries	Effect
A	397,472	237,006	0.6
B	48,347,395	2,973,118	0.1
C	46,080,139	1,376,499,745,600	29,871.9
E	19,208,299	290,999,305,488	15,149.7
F	39,942,638	2,844,513,432,416,357,974,016	71,214,961,626,128.9
G	561,202	912,194	1.6
I	95,730,831	116,387,431,997,281,136	1,215,777,934.7
J	24,619,563	15,424,761,908	626.5
K	250,971,828	2,544,758,557,238,438	10,139,618.4

"Misses" is again the number of cache misses in our BDD-based implementation, and "Table entries" is the total number of table entries in the naive implementation. To roughly estimate the effect of the BDD-representation, we conservatively assume that each table entry results in just a single cache miss; thus, "Effect" compares "Table entries" to "Misses". The few instances where the effect is less than one correctly identify benchmark formulas where the BDDs are less necessary, but are also artifacts of our conservative assumption. Conversely, the extremely high effects are associated with formulas that could not possibly be decided without BDDs. Of course, the use of BDD-structures completely dominates all other optimizations, since no implementation could realistically be based on the naive table representation.

The BDD-representation was the first breakthrough of the MONA implementation, and the other "secrets" should really be viewed with this as baseline. The first implementation did not actually use tables but a conjunctive normal form. Nevertheless, the effect of switching to BDDs was stunning.

4.5 DAGification

Internally, MONA is divided into a front-end and a back-end. The front-end parses the input and builds a data structure representing the automata-theoretic operations that will calculate the resulting automaton. The back-end then inductively carries out these operations.

The generated data structure is often seen to contain many common subformulas. This is particularly true when they are compared relative to *signature equivalence*, which holds for two formulas ϕ and ϕ' if there is an order-preserving renaming of the variables in ϕ (increasing with respect to the indices of the variables) such that the representations of ϕ and ϕ' become identical.

A property of the BDD representation is that the automata corresponding to signature-equivalent trees are isomorphic in the sense that only the node indices differ. This means that intermediate results can be reused by simple exchanges of node indices. For this reason, MONA represents the formulas in a DAG (Directed Acyclic Graph), not a tree. The DAG is conceptually constructed from the tree using a bottom-up collapsing process, based on the signature equivalence relation as described in [11].

Clearly, constructing the DAG instead of the tree incurs some overhead, but the following experiments show that the benefits are significantly larger:

Benchmark	Nodes		Time		Effect
	Tree	DAG	Tree	DAG	
A	2,532	296	1.7 sec	0.4 sec	4.3
B	873	259	79.2 sec	33.1 sec	2.4
C	5,432	461	40.1 sec	15.1 sec	2.7
D	3,038	270	∞	137.1 sec	∞
E	4,560	505	20.5 sec	8.0 sec	2.6
F	1,997	505	49.1 sec	13.2 sec	3.7
G	56,932	1,199	∞	3.2 sec	∞
H	8,180	743	∞	30.4 sec	∞
I	14,058	1,396	107.1 sec	40.3 sec	2.7
J	278,116	6,314	∞	20.6 sec	∞
K	777	273	284.0 sec	139.9 sec	2.0
L	1,504	388	∞	136.4 sec	∞

"Nodes" shows the number of nodes in the representation of the formula. "Tree" is the number of nodes using an explicit tree representation, while "DAG" is the number of nodes after DAGification. "Time" shows the running times for the same two cases. "Effect" shows the "Tree" running time compared to the "DAG" running time. The DAGification is seen to provide a substantial and often essential gain in efficiency.

The effects reported sometimes benefit from the fact that the restriction technique presented in the following subsection knowingly generates redundant formulas. This explains some of the failures observed.

4.6 Three-valued logic and automata

In earlier versions of Mona, we struggled with the issue of encoding first-order variables as second-order variables—that's the standard technique for monadic second-order logics, but it raises the issue of *restrictions*: the common phenomenon that a formula ϕ makes sense, relative to some exterior conditions, only when an associated restriction holds. The restriction is also a formula, and the main issue is that ϕ is now essentially undefined outside the restriction. Later, when we chose to base Mona on WS1S instead of a monadic second-order logic on strings, we sometimes encountered state space explosions when we constrained variables in order to emulate the string-based semantics.

The nature of these problems is very technical, but fortunately they can be solved through a theory of restriction couched in a three-valued logic [19]. Under this view, a *restricted subformula* ϕ is associated with a restriction ϕ_R different from *true*; an *unrestricted formula* is associated with a restriction ϕ_R that is *true*. We do not outline the theory of restrictions here, except for noting that restriction of a conjunction (or disjunction) is the conjunction of the restrictions of the conjuncts (or disjuncts).

According to [19], we can guarantee that the WS1S framework handles all formulas written in the earlier string logic, even with intermediate automata that are no bigger than when run through the original decision procedure. Also, the running time of the original procedure may be asymptotically worse than with the WS1S formulation. Unfortunately, there is no way of disabling this feature to provide a quantitative comparison.

4.7 Formula reductions

Formula reduction is a means of "optimizing" the formulas in the DAG before translating them into automata. The reductions are based on a syntactic analysis that attempts to identify valid subformulas and equivalences among subformulas.

There are some non-obvious choices here. How should computation resources be apportioned to the reduction phase and to the automata calculation phase? Must reductions guarantee that automata calculations become faster? Should the two phases interact? Our answers are based on some trial and error along with some provisions to cope with subtle interactions with other of our optimization secrets.

Mona 1.4 performs three kinds of formula reductions: 1) simple equality and boolean reductions, 2) special quantifier reductions, and 3) special conjunction reductions. The first kind can be described by simple rewrite rules (only some typical ones are shown):

$$
\begin{aligned}
x = x &\rightsquigarrow \mathit{true} & \qquad \phi \wedge \phi &\rightsquigarrow \phi \\
\mathit{true} \wedge \phi &\rightsquigarrow \phi & \neg\neg\phi &\rightsquigarrow \phi \\
\mathit{false} \wedge \phi &\rightsquigarrow \mathit{false} & \neg\mathit{false} &\rightsquigarrow \mathit{true}
\end{aligned}
$$

These rewrite steps are guaranteed to reduce complexity, but will not cause significant improvements in running time, since they all either deal with constant size automata or rarely apply in realistic situations. Nevertheless, they are extremely cheap, and they may yield small improvements, in particular on machine generated MONA code.

The second kind of reductions can potentially cause tremendous improvements. The non-elementary complexity of the decision procedure is caused by the automaton projection operations, which stem from quantifiers. The accompanying determinization construction may cause an exponential blow-up in automaton size. Our basic idea is to apply a rewrite step resembling *let*-reduction, which removes quantifiers:

$\exists X : \phi \rightsquigarrow \phi[T/X]$ provided that $\phi \Rightarrow X = T$ is valid, and T is some term satisfying $FV(T) \subseteq FV(\phi)$

where $FV(\cdot)$ denotes the set of free variables. For several reasons, this is not the way to proceed in practice. First of all, finding terms T satisfying the side condition can be an expensive task, in worst case non-elementary. Secondly, the translation into automata requires the formulas to be "flattened" by introduction of quantifiers such that there are no nested terms. So, if the substitution $\phi[T/X]$ generates nested terms, then the removed quantifier is recreated by the translation. Thirdly, when the rewrite rule applies in practice, ϕ usually has a particular structure as reflected in the following more restrictive rewrite rule chosen in MONA:

$\exists X : \phi \rightsquigarrow \phi[Y/X]$ provided that $\phi \equiv \cdots \wedge X = Y \wedge \cdots$, and Y is some variable other than X

In contrast to equality and boolean reductions, this rule is not guaranteed to improve performance, since substitutions may cause the DAG reuse degree to decrease.

The third kind of reductions applies to conjunctions, of which there are two special sources. One is the formula flattening just mentioned; the other is the formula restriction technique mentioned in Section 4.6. Both typically introduce many new conjunctions. Studies of a graphical representation of the formula DAGs (MONA can create such graphs automatically) led us to believe that many of these new conjunctions are redundant. A typical rewrite rule addressing such redundant conjunctions is the following:

$\phi_1 \wedge \phi_2 \rightsquigarrow \phi_1$ provided that $\mathit{unrestr}(\phi_2) \subseteq \mathit{unrestr}(\phi_1) \cup \mathit{restr}(\phi_1)$ and $\mathit{restr}(\phi_2) \subseteq \mathit{restr}(\phi_1)$

Here, $\mathit{unrestr}(\phi)$ is the set of unrestricted conjuncts in ϕ, and $\mathit{restr}(\phi)$ is the set of restricted conjuncts in ϕ. This reduction states that it is sufficient to assert ϕ_1 when $\phi_1 \wedge \phi_2$ was originally asserted in situations where the unrestricted conjuncts of ϕ_2 are already conjuncts of ϕ_1—whether restricted or not—and the restricted conjuncts of ϕ_2 are unrestricted conjuncts of ϕ_1. It is not sufficient that they be restricted conjuncts of ϕ_1, since the restrictions may not be the same in ϕ_1.

With the DAG representation of formulas, the reductions just described can be implemented relatively easily in MONA. The table below shows the effects of performing the reductions on the benchmark formulas:

Benchmark	Hits			Time					Effect
	Simple	Quant.	Conj.	None	Simple	Quant.	Conj.	All	
A	12	8	22	0.8 sec	0.7 sec	0.7 sec	0.7 sec	0.4 sec	2.0
B	10	45	0	58.2 sec	58.8 sec	56.2 sec	56.8 sec	33.1 sec	1.8
C	9	13	8	43.7 sec	41.9 sec	37.1 sec	42.9 sec	15.1 sec	2.9
D	4	28	27	542.7 sec	536.1 sec	296.0 sec	404.7 sec	137.1 sec	4.0
E	5	6	19	22.6 sec	23.4 sec	16.6 sec	22.7 sec	8.0 sec	2.8
F	3	1	1	28.3 sec	29.9 sec	27.0 sec	27.2 sec	13.2 sec	2.1
G	65	318	191	6.1 sec	5.9 sec	6.1 sec	5.9 sec	3.2 sec	1.9
H	35	32	81	104.1 sec	102.6 sec	71.0 sec	98.5 sec	30.4 sec	3.4
I	102	218	7	76.2 sec	76.5 sec	75.0 sec	76.0 sec	40.3 sec	1.9
J	91	0	1	37.3 sec	37.9 sec	37.6 sec	37.0 sec	20.6 sec	1.9
K	9	4	1	313.7 sec	267.9 sec	240.3 sec	302.6 sec	139.9 sec	2.3
L	4	4	18	∞	∞	∞	∞	136.4 sec	∞

"Hits" shows the number of times each of the three kinds of reduction is performed; "Time" shows the total running time in the cases where no reductions are performed, only the first kind of reductions are performed, only the second, only the third, and all of them together. "Effect" shows the "None" times compared to the "All" times. All benchmarks gain from formula reductions, and in a single example this technique is even necessary. Note that most often all three kinds of reductions must act in unison to obtain significant effects.

A general benefit from formula reductions is that tools generating Mona formulas from other formalisms may generate naive and voluminous output while leaving optimizations to Mona. In particular, tools may use existential quantifiers to bind terms to fresh variables, knowing that Mona will take care of the required optimization.

5 Future Developments

Several of the techniques described in the previous section can be further refined of course. The most promising ideas seem however to concentrate on the BDD representation. In the following, we describe three such ideas.

It is a well-known fact [6] that the ordering of variables in the BDD automata representation has a strong influence on the number of BDD nodes required. The impact of choosing a good ordering can be an exponential improvement in running times. Finding the optimal ordering is an NP-complete problem, but we plan to experiment with the heuristics that have been suggested [9].

We have sometimes been asked: "Why don't you encode the states of the automata in BDDs, since that is a central technique in model checking?". The reason is very clear: there is no obvious structure to the state space in most cases that would lend itself towards an efficient BDD representation. For example, consider the consequences of a subset construction or a minimization construction, where similar states are collapsed; in either case, it is not obvious how to represent the new state. However, the ideas are worth investigating.

For our tree automata, we have experimentally observed that the use of guides produce a large number of component automata many of which are almost identical. We will study how to compress this representation using a BDD-like global structure.

6 Conclusion

The presented techniques reflect a lengthy Darwinian development process of the Mona tool in which only robust and useful ideas have survived. We have not mentioned here the many ideas that failed or were surpassed by other techniques. Our experiences confirm the maxim that optimizations must be carried out at all levels and that no single silver bullet is sufficient. We are confident that further improvements are still possible.

Many people have contributed to earlier versions of Mona, in particular we are grateful to David Basin, Morten Biehl, Jacob Elgaard, Jesper Gulmann, Jacob Jensen, Michael Jørgensen, Bob Paige, Theis Rauhe, and Anders Sandholm. We also thank the Mona users who kindly provided the benchmark formulas.

References

[1] Parosh Aziz Abdulla, Ahmed Bouajjani, Bengt Jonsson, and Marcus Nilsson. Handling global conditions in parameterized system verification. In *Computer Aided Verification, CAV '99*, volume 1633 of *LNCS*, 1999.

[2] David Basin and Nils Klarlund. Hardware verification using monadic second-order logic. In *Computer Aided Verification, CAV '95*, volume 939 of *LNCS*, 1995.

[3] David Basin and Nils Klarlund. Automata based symbolic reasoning in hardware verification. *Formal Methods in Systems Design*, 13(3):255–288, 1998.

[4] Morten Biehl, Nils Klarlund, and Theis Rauhe. Algorithms for guided tree automata. In *WIA '96*, volume 1260 of *LNCS*, 1996.

[5] Jean-Paul Bodeveix and Mamoun Filali. FMona: a tool for expressing validation techniques over infinite state systems. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2000*, volume 1785 of *LNCS*, 2000.

[6] R.E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, Aug 1986.

[7] J.R. Büchi. Weak second-order arithmetic and finite automata. *Z. Math. Logik Grundl. Math.*, 6:66–92, 1960.

[8] J.R. Büchi. On a decision method in restricted second-order arithmetic. In *Proc. Internat. Cong. on Logic, Methodol., and Philos. of Sci.* Stanford University Press, 1962.

[9] K. M. Butler, D. E. Ross, R. Kapur, and M. R. Mercer. Heuristics to compute variable orderings for efficient manipulation of ordered binary decision diagrams. In *Proc. ACM/IEEE Design Automation Conference (DAC)*, pages 417–420, 1991.

[10] Niels Damgaard, Nils Klarlund, and Michael I. Schwartzbach. YakYak: Parsing with logical side constraints. In *Proceedings of DLT'99*, 1999.

[11] Jacob Elgaard, Nils Klarlund, and Anders Møller. Mona 1.x: New techniques for WS1S and WS2S. In *Computer Aided Verification, CAV '98*, volume 1427 of *LNCS*, 1998.

[12] Jacob Elgaard, Anders Møller, and Michael I. Schwartzbach. Compile-time debugging of C programs working on trees. In *Proceedings of European Symposium on Programming Languages and Systems*, volume 1782 of *LNCS*, 2000.

[13] C.C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.

[14] J.G. Henriksen, J. Jensen, M. Jørgensen, N. Klarlund, B. Paige, T. Rauhe, and A. Sandholm. Mona: Monadic second-order logic in practice. In *Tools and Algorithms for the Construction and Analysis of Systems, First International Workshop, TACAS '95*, volume 1019 of *LNCS*, 1996.

[15] Thomas Hune and Anders Sandholm. A case study on using automata in control synthesis. In *Proceedings of FASE'00*, 2000.

[16] J.L. Jensen, M.E. Jørgensen, N. Klarlund, and M.I. Schwartzbach. Automatic verification of pointer programs using monadic second-order logic. In *SIGPLAN '97 Conference on Programming Language Design and Implementation,*, pages 226–234, 1997.

[17] N. Klarlund, M. Nielsen, and K. Sunesen. A case study in automated verification based on trace abstractions. In M. Broy, S. Merz, and K. Spies, editors, *Formal System Specification, The RPC-Memory Specification Case Study*, volume 1169 of *LNCS*, pages 341–374. Springer Verlag, 1996.

[18] N. Klarlund and T. Rauhe. BDD algorithms and cache misses. Technical report, BRICS Report Series RS-96-5, Department of Computer Science, University of Aarhus, 1996.

[19] Nils Klarlund. A theory of restrictions for logics and automata. In *Computer Aided Verification, CAV '99*, volume 1633 of *LNCS*, 1999.

[20] Nils Klarlund and Anders Møller. *MONA Version 1.3 User Manual.* BRICS Notes Series NS-98-3 (2.revision), Department of Computer Science, University of Aarhus, October 1998. URL: `http://www.brics.dk/mona/manual.html`.

[21] A.R. Meyer. Weak monadic second-order theory of successor is not elementary recursive. In R. Parikh, editor, *Logic Colloquium, (Proc. Symposium on Logic, Boston, 1972)*, volume 453 of *LNCS*, pages 132–154, 1975.

[22] Anders Møller. MONA homepage. URL: `http://www.brics.dk/mona/`.

[23] Frank Morawietz and Tom Cornell. The logic-automaton connection in linguistics. In *Proceedings of LACL 1997*, volume 1582 of *LNAI*, 1997.

[24] Paritosh K. Pandya. DCVALID 1.2. Technical report, Theoretical Computer Science Group, Tata Institute of Fundamental Research, 1997.

[25] Anders Sandholm and Michael I. Schwartzbach. Distributed safety controllers for web services. In *FASE*, volume 1382 of *LNCS*, 1998.

[26] Thomas R. Shiple, James H. Kukula, and Rajeev K. Ranjan. A comparison of Presburger engines for EFSM reachability. In *CAV*, volume 1427 of *LNCS*, 1998.

[27] Mark A. Smith and Nils Klarlund. Verification of a sliding window protocol using IOA and Mona, 1999. Submitted for publication.

[28] W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–191. MIT Press/Elsevier, 1990.

[29] S. Yang. Logic synthesis and optimization benchmarks user guide version 3.0. In *Tech. Rep. MCNC*, 1991.