



Efficient Simplification of Bisimulation Formulas

Uffe H. Engberg¹
Kim S. Larsen^{2,3}

ABSTRACT The problem of checking or optimally simplifying bisimulation formulas is likely to be computationally very hard. We take a different view at the problem: we set out to define a very fast algorithm, and then see what we can obtain. Sometimes our algorithm can simplify a formula perfectly, sometimes it cannot. However, the algorithm is extremely fast and can, therefore, be added to formula-based bisimulation model checkers at practically no cost. When the formula can be simplified by our algorithm, this can have a dramatic positive effect on the better, but also more time consuming, theorem provers which will finish the job.

1 Introduction

The need for validity checking or optimal simplification of first order bisimulation formulas has arisen from recent work on symbolic bisimulation checking of *value-passing calculi* [4, 9, 15]. The NP-completeness of checking satisfiability of propositional formulas [3] implies that validity checking of that class of formulas is co-NP complete. Additionally, checking of quantified formulas is P-space hard [7], so there is not much hope for a fast algorithm for deciding exactly when a bisimulation formula is valid.

Instead, we set out to solve the problem of what you can get for free, i.e., to what extent is it possible to decide validity simply while reading the formula? As it turns out, there is almost nothing that can be done in linear time. The most simple tasks of storing and retrieving information about variables will cost $O(n \log n)$. So, we allowed ourselves this extra log-factor and changed the question to what you can get *almost* for free. As we shall demonstrate in this paper, the algorithm we have designed is very fast. Not alone does it run in $O(n \log n)$; the constant is also very small. On average, we read through the formulas at a rate of about 75 Kbytes per second. Of course, this is only interesting if the algorithm outputs

¹BRICS (Basic Research in Computer Science, a Centre of the Danish National Research Foundation), Department of Computer Science, University of Aarhus, Denmark.

²Department of Mathematics and Computer Science, Odense University, Denmark.

³The initial part of this work was done while this author was at University of Aarhus.

useful answers reasonably frequently, i.e., in the absence of an obvious notion of optimal simplification (to a minimal equivalent formula), if the algorithm can reasonably frequently guarantee that the formula is valid, rule out that the formula could be valid, or maybe simplify a huge formula to a much smaller equivalent one. It is not easy to measure how often the algorithm produces a useful answer, but through examples, we show that there are families of process expressions which give rise to formulas, where our algorithm is successful.

2 Applications

The algorithm makes a single pass over the formula making no assumptions about the variable names. Notice that this also implies that if formulas are passed on to our algorithm from another program, they do not have to be saved at any point, but can be passed on to our program via pipelining.

As already mentioned, in addition to validity checking, we can also simplify formulas. This greatly increases the usefulness of this work. Even if our algorithm fails to prove the validity of a formula, it will quite often simplify the formula so drastically that the validity of the formula (or the opposite) can easily be asserted by the user. Another possibility is to check the simplified formula for validity using other tools. Tools that would succeed more often, but the complexity of which would make it impossible to work on the original formula. The main advantage of simplifying formulas is that the algorithm can be built into the program generating the bisimulation formulas and simplify the formula on the fly. Also, intermediate simplifications can be used to prune the generation of other subformulas. This can be of vital importance since the formulas can become exponentially large.

Our algorithms work just as well on formulas with free variables as on closed formulas. This means that by simplifying a formula with free variables, we actually characterize the conditions, in terms of the free variables, under which two processes are bisimilar.

Though this work is motivated by the need for checking bisimulation formulas, the method is sound for first order formulas in general, so there may be other interesting applications. Depending on the concrete application, extra facilities can be built in at no extra cost (with regards to asymptotic time complexity). For example, when our algorithm has deduced that some variables must be equivalent, in the rest of the formula, all occurrences of these variables can be replaced by one representative.

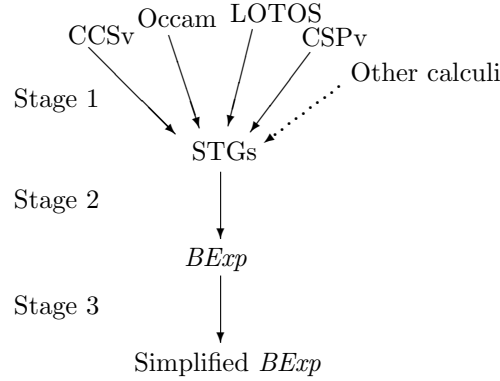


FIGURE 1. Verification stages.

3 Bisimulation Formulas

In this section, we rather briefly describe the context in which bisimulation formulas arise and how they are constructed. This section can be skipped if one believes that such formulas characterizing bisimilarity of two processes can be generated.

In [4, 9], bisimulation formulas are used in a three stage process (cf. figure 1) of verifying bisimilarity of value-passing programs ([15] obtain similar formulas, albeit with a different approach).

However, because we mainly want to convey the intuition of how bisimulation formulas are obtained we chose a simple presentation avoiding the notions of symbolic bisimulation and symbolic operational semantics in [9].

Stage 1

In the first stage, symbolic transition graphs (a generalization of the standard notion of labelled transition graphs) are generated from terms of some value-passing language, say the full CCS calculus [14], CCSv, with terms

$$t ::= \mathbf{0} \mid \tau.t \mid a!e.t \mid a?x.t \mid \mathbf{if} \ c \ \mathbf{then} \ t \mid t + t \mid t | t \mid t \setminus b \mid A(\bar{e})$$

where a and b range over channel names, x range over a set, V , of variables, e range over a given set of expressions (including V) over some value domain D with elements d , c is a condition ranging over boolean expressions, and A range over names with associated declaration of the form $A(\bar{x}) \leftarrow t$. As usual the set of programs or processes, $Proc$, are the closed terms (no free variables, where $?x$ acts as binder).

Following [14], $\mathbf{0}$ is the process having no actions, whereas the prefixed process $\tau.p$ can make an internal action and then act as p . Similarly, there are input (output) prefixes $a?x$. ($a!x$.) for receiving (sending) values on

$$\begin{array}{ll}
\tau.p \xrightarrow{\tau} p, & a!d.p \xrightarrow{a!d} p, & a?x.p \xrightarrow{a?x} p \\
\mathbf{if} \top \mathbf{then} p \xrightarrow{\alpha} p' & \text{if } p \xrightarrow{\alpha} p' \\
p + q \xrightarrow{\alpha} p' & \text{if } p \xrightarrow{\alpha} p' \\
p \mid q \xrightarrow{\alpha} p' \mid q & \text{if } p \xrightarrow{\alpha} p' \\
p \mid q \xrightarrow{\tau} p'[d/x] \mid q' & \text{if } p \xrightarrow{a?x} p', q \xrightarrow{a!d} q' \\
p \setminus b \xrightarrow{\alpha} p' \setminus b & \text{if } p \xrightarrow{\alpha} p', b \notin \alpha \\
A(\bar{d}) \xrightarrow{\alpha} p' & \text{if } t[\bar{d}/\bar{x}] \xrightarrow{\alpha} p' \text{ and } A(\bar{x}) \Leftarrow t
\end{array}$$

FIGURE 2. Operational interpretation of CCSv.

channel a . The expression **if** c **then** p can do the actions of p , provided the condition c evaluates to true, \top . The process $p + p'$ acts as either p or p' . In $p \mid q$, the parallel subprocesses p and q are (continuously) acting in parallel through either an internal communication of a value or a separate action of one of the subprocesses. $p \setminus b$ can do those external actions of p not involving communication on b . Finally, $A(\bar{e})$ acts as the process obtained by instantiating the associated term with the parameters of A .

Formally processes are equipped with action relations,

$$\xrightarrow{\tau}, \xrightarrow{a!d} \subseteq Proc \times Proc, \quad \xrightarrow{a?x} \subseteq Proc \times (D \rightarrow Proc), \quad \text{where}$$

$p \xrightarrow{a?x} \lambda d.p'[d/x]$ if $p \xrightarrow{a?x} p'$ and the remaining are defined as the least relations satisfying the rules of figure 2 (symmetric rules left out) where α range over actions of the form τ , $a!d$ or $a?x$, and $b \notin \alpha$ means $b \neq a$.

Bisimilarity of processes is defined by means of bisimulations. A relation, \mathcal{R} , on processes is a bisimulation, if it is symmetric and $(p, q) \in \mathcal{R}$ implies

1. $p \xrightarrow{\tau} p' \Rightarrow \exists q': q \xrightarrow{\tau} q' \text{ and } (p', q') \in \mathcal{R}$
2. $p \xrightarrow{a!d} p' \Rightarrow \exists q': q \xrightarrow{a!d} q' \text{ and } (p', q') \in \mathcal{R}$
3. $p \xrightarrow{a?x} p' \Rightarrow \exists q': q \xrightarrow{a?x} q', \forall d \in D: (p'd, q'd) \in \mathcal{R}$

Then p and q are bisimilar, $p \sim q$, iff $p \mathcal{R} q$ for some bisimulation \mathcal{R} .

This is in fact what is usually called *late* bisimulation equivalence for value-passing processes. The *early* version (used in [14]) is obtained by merely moving the universal quantification completely to left of the right hand side of the implication in clause 3—the effect being that a (possibly new) q' can be found for each d .

Stage 1 and 2 of the verification process (and the associated theory) can equally well be carried through for early bisimulation. The simplification in stage 3, which is the main concern of the article, is essentially independent of the bisimulation equivalence in question. Henceforth we shall tacitly stick to late bisimulation equivalence.

$$\begin{array}{l}
\tau.t \xrightarrow{\top, \tau} t, \quad a!e.t \xrightarrow{\top, a!e} t, \quad a?x.t \xrightarrow{\top, a?\#} t \\
\mathbf{if } c \mathbf{ then } t \xrightarrow{c \wedge c', \alpha} p' \quad \text{if } t \xrightarrow{c', \alpha} t' \\
t + u \xrightarrow{c, \alpha} t' \quad \text{if } t \xrightarrow{c, \alpha} t' \\
t | u \xrightarrow{c, \alpha} t' | u \quad \text{if } t \xrightarrow{c, \alpha} t' \\
t | u \xrightarrow{c \wedge c', \tau} t'[e/\#] | u' \quad \text{if } t \xrightarrow{c, a?\#} t', u \xrightarrow{c', a!e} u' \\
t \setminus b \xrightarrow{c, \alpha} t' \setminus b \quad \text{if } t \xrightarrow{c, \alpha} t', b \notin \alpha \\
A(\bar{e}) \xrightarrow{c, \alpha} t' \quad \text{if } t[\bar{e}/\bar{x}] \xrightarrow{c, \alpha} t' \text{ and } A(\bar{x}) \Leftarrow t
\end{array}$$

FIGURE 3. Generating symbolic transition graphs from CCSv.

We shall now generalize the notion of bisimilarity to open terms, i.e., parametrized programs, in such a way that it allows us to formally characterize the conditions, in terms of the parameters, under which the (parametrized) programs are bisimilar.

An *environment*, is a function from V to D , typically denoted ρ . It is then standard how to define when ρ *satisfies* a formula c , written $\rho \models c$. Given an environment ρ and a term t , we use t_ρ to denote the closed term obtained by substituting $\rho(x)$ for each free variable x in t . \sim is then generalized by

$$t \sim^c u \text{ iff } \forall \rho \models c: t_\rho \sim u_\rho.$$

A boolean expression, c , is *most general* for terms t and u if it satisfies $t \sim^c u$ and if $t \sim^{c'}$ u , then c' implies c .

Now for *symbolic transition graphs*. A symbolic transition graph is merely a labelled directed graph satisfying certain well-formedness conditions. Nodes are labelled with variables and the edges (\mapsto) labelled with (symbolic) actions of the form τ , $a!e$ or $a?x$, each guarded with some boolean condition. Intuitively, a graph is well-formed if the (free) variables of the expressions labelling any edge leading from node n to node m only mention variables of n and a variable, x , of m not in n is due to the edge being labelled with an input action $a?x$ binding x .

The advantage of symbolic transition graphs is their independence of a particular process language. Various calculi can be translated into them as we shall see now for CCSv.

It is relatively simple to *generate a symbolic transition graph* from a CCSv term. Edges are (modulo symbolic actions) inferred similarly to the action relations in figure 2, except that conditions met are collected (taking the conjunction) and used as guarding condition. For example, consider the rules of figure 3 (where $\#$ is a special place holder variable not in V) and let $t \xrightarrow{c, a?x} t''$ if $t \xrightarrow{c, a?\#} t'$, $t'' = t'[x/\#]$ and x is the first variable not free in t' . A (not particularly sophisticated) translation is then obtained by letting nodes being terms labelled with their free variables and edges

$$\begin{array}{ll}
m_\rho \xrightarrow{\tau} n_\rho & \text{if } \rho \models c \text{ and } m \xrightarrow{c, \tau} n \\
m_\rho \xrightarrow{a! \rho(e)} n_\rho & \text{if } \rho \models c \text{ and } m \xrightarrow{c, a!e} n \\
m_\rho \xrightarrow{a?} \lambda d. n_{\rho[x \mapsto d]} & \text{if } \rho \models c \text{ and } m \xrightarrow{c, a?x} n
\end{array}$$

FIGURE 4. Operational interpretation of symbolic transition graphs.

given according to the rules. To get finite graphs more often, one can apply modifications of standard techniques used to avoid state explosion in tools for process languages without value-passing.

Below, two processes are shown together with their associated graphs.

$$A_1(x, y) \stackrel{\text{def}}{=} \tau. \mathbf{0} + \tau. \text{if } x \equiv y \text{ then } \tau. \mathbf{0} \quad A_3(x, y) \stackrel{\text{def}}{=} \tau. (\text{if } x \equiv y \text{ then } \tau. \mathbf{0} \\
+ \text{if } x \equiv y \text{ then } \tau. \mathbf{0})$$



To shorten the presentation, all nodes except for root nodes, have been omitted from the graphs and in the conditions, conjunctions with \top have been removed. The initial conditions of the graph g_1 associated with A_1 are \top , whereas the last condition is $x \equiv y$ (or rather $x \equiv y \wedge \top$), i.e., x and y should have equivalent values.

The symbolic transition graphs can now be given an operational interpretation which closely relates the interpretation of translated CCSv terms with the original interpretation of the terms. Given an environment ρ and a node n the associated state is denoted n_ρ . The transition relations are then given by the rules of figure 4, where $\rho[x \mapsto d]$ denotes an environment identical to ρ except that it yields d for x . Notice that $\lambda d. n_{\rho[x \mapsto d]}$ applied to a $d' \in D$ gives the state $n_{\rho[x \mapsto d']}$.

Two graphs are *bisimilar* [4, 9, 10] iff the two terms they are translated from, are bisimilar in the traditional sense.

Stage 2

In the next stage, an algorithm is used for finding a first order boolean expression *mgb*, called the *most general boolean*; a bisimulation formula characterizing the conditions for which two finite symbolic transition graphs are bisimilar.

Intuitively, two processes are bisimilar if, whenever one process can do an action, the other has a matching action such that the resulting two processes are again bisimilar. This is reflected in the bisimulation formula. For example, the fact that g_3 must match (see figure 5) an action corresponding




FIGURE 5. Generating symbolic bisimulation formulas.

to the left edge of g_1 , is captured in the mgb_{g_1, g_3} subformula

$$\mathbb{T} \rightarrow \bigvee \frac{\mathbb{T} \wedge [\mathbb{T} \wedge (x \equiv y \rightarrow \mathbb{F})]}{x \equiv y \wedge [\mathbb{T} \wedge \mathbb{T}]}$$

If an instantiation of variables satisfies the guarding condition \mathbb{T} of g_1 , then it must satisfy a guarding condition, \mathbb{T} or $x \equiv y$, of g_3 as well as the mgb , $[\mathbb{T} \wedge (a \equiv y \rightarrow \mathbb{F})]$ or $[\mathbb{T} \wedge \mathbb{T}]$, of the corresponding nodes.

More precisely, for a pair of nodes m and m' the bisimulation formula $mgb_{m, m'}$ is constructed recursively as a large conjunction of subformulas each capturing the possibility of matching an m -action with m' -actions (and vice versa) and recursively being able to do so.

If there are no actions to match, the subformula is simply \mathbb{T} . Otherwise, for each action $m \xrightarrow{c, \alpha} n$, the subformula expresses that the guarding condition c imply the possibility of finding a matching m' -action.

Hence, if there are no possible m' -action of the same form, communicating on the same channel, the subformula is merely $c \rightarrow \mathbb{F}$. Otherwise, the constructed subformula depends on the form of α .

If α is an output action, $a!e$, the subformula is

$$c \rightarrow \bigvee_i [c'_i \wedge e = e'_i \wedge mgb_{n, n'_i}], \text{ where } m' \xrightarrow{c'_i, a!e'_i} n'_i$$

meaning that whenever a process environment satisfies c there is at least one m -action which in the same environment and on the same channel can send the same value.

Similarly, for input, universal quantification is used to express that for all values received, the processes are bisimilar, i.e., for α of the form $a?x$

$$c \rightarrow \bigvee_i [c'_i \wedge \forall x, x'_i: x = x'_i \rightarrow mgb_{n, n'_i}]$$

For internal actions, the subformula is just $c \rightarrow \bigvee_i [c'_i \wedge mgb_{n, n'_i}]$.

The construction clearly stops if there are no cycles in the graphs. But if there are, it will go on forever. However, as we proceed down in the subformulas, more and more constraints are accumulated through the conditions. Hence, intuitively, if the accumulated constraints imply the mgb we are about to construct, the recursion can in principle be stopped with subformula \mathbb{T} . It turns out that for the examples considered here, this is in fact the case whenever previously visited pairs of nodes are met again.

Stage 3

In the final stage, validity of the bisimulation formula is checked. If it is valid, the original programs are bisimilar under all instantiations. Otherwise, the formula expresses the weakest conditions on the instantiations for which they are bisimilar. Later, in section 9, we shall see that mgb_{g_1, g_3} is in fact valid.

4 Logic Notions

We now fix the set of bisimulation formulas and introduce a few general notions from formal logic.

Formally, the class of formulas, which we will work with in the rest of this paper, is defined by the syntax

$$E ::= P \mid E \wedge E \mid E \vee E \mid P \rightarrow E \mid \forall x: E \quad \text{and} \quad P ::= \top \mid \text{F} \mid x \equiv y$$

As usual, formulas are closed if they have no free variables. Notice that bisimulation formulas only have universal quantification. The binary predicate symbol \equiv is assumed to be interpreted as an equivalence relation \equiv_D over a nonempty domain D . An environment satisfies a set of formulas Γ , if it satisfies each formula of Γ . Γ *semantically entails* E , written $\Gamma \models E$, if E is satisfied by any environment satisfying Γ . E is *valid*, $\models E$, if $\emptyset \models E$.

The class of bisimulation formulas is a subset of the class of all quantified formulas and checking validity of bisimulation formulas is not P-space hard. An easy reduction shows that checking satisfiability with respect to an environment is NP complete, so presumably checking validity of bisimulation formulas is co-NP complete.

We now state some general properties of entailment relevant for the development of the algorithm. For simplicity, we write Γ as E_1, \dots, E_n when $\Gamma = \{E_1, \dots, E_n\}$. Similarly, we write Γ, Δ for $\Gamma \cup \Delta$.

Theorem 1 (Entailment)

- a) If $\Gamma \models E$ then $\Gamma, \Delta \models E$ (Ext)
- b) If $E \in \Gamma$ then $\Gamma \models E$ (Rep)
- c) If $\Gamma \models E$ and $\Gamma, E \models E'$ then $\Gamma \models E'$ (Cut)
- d) If $\Gamma \models E$ and $E \models E'$ then $\Gamma \models E'$ (Trans)
- e) $\Gamma \models E$ and $\Gamma \models E'$ iff $\Gamma \models E \wedge E'$ (Conj)
- f) $\Gamma, E \models E'$ iff $\Gamma \models E \rightarrow E'$ (Imp)

Proof Standard, see [16], for example. □

In general, there is not any similar disjunction theorem allowing both introduction and elimination to the right. However, from the entailment theorems and a few tautologies, we get proposition 2. As a consequence of \equiv_D being an equivalence relation we also have proposition 3.

Proposition 2 If $\Gamma \models E$ or $\Gamma \models E'$, then $\Gamma \models E \vee E'$.

Proposition 3 a) $\models x \equiv x$ b) $x \equiv y \models y \equiv x$ c) $x \equiv y, y \equiv z \models x \equiv z$

5 The Abstract Algorithm

We now set out to design an algorithm for checking validity of formulas with an equivalence predicate. We keep it as abstract as possible to allow for a large degree of freedom in the choice of data structures in the actual implementation.

Intuitively, the idea of the algorithm is to collect in a relation R (over variables) information about variables known to be equivalent when checking subformulas. For instance, checking the validity of a formula like $x \equiv y \rightarrow E$ is reduced to checking E under the assumption that x and y are equivalent, i.e., (x, y) is added to R and E then checked. To exploit that \equiv is an equivalence relation, the symmetric and transitive closure, is taken before proceeding to E . However, when checking E of the formula $\forall x : E$ the situation is quite the opposite. Since a new scope is entered, all previous collected information in R concerning x , must be removed before E is checked.

Formally, for R denote the symmetric and transitive closure by R^\oplus , the reflexive closure, $R \cup \{(x, x) \mid x \in V\}$, by R^0 and the removal of x , $\{(y, z) \in R \mid y \neq x, z \neq x\}$, by $R \setminus x$. Notice, $R \setminus x \subseteq R$, and if R is symmetrically and transitively closed, then so is $R \setminus x$.

In order to connect with the logic, we associate with R the set of formulas $R_\equiv \stackrel{\text{def}}{=} \{x \equiv y \mid (x, y) \in R\}$. The notions of closures and removal extend to R_\equiv in the natural way: $R_\equiv \setminus x$ is $(R \setminus x)_\equiv$ etc.

The algorithm is conveniently described using Kleene's three-valued logic [11], the three truth-values being **t** for "true", **f** for "false" and **u** for "undefined"/ "unknown". The Kleene truth tables for conjunction, \wedge_K , disjunction, \vee_K , and implication, \rightarrow_K , are:

| \wedge_K | t | f | u |
|------------|----------|----------|----------|
| t | t | f | u |
| f | f | f | f |
| u | u | f | u |

| \vee_K | t | f | u |
|----------|----------|----------|----------|
| t | t | t | t |
| f | t | f | u |
| u | t | u | u |

| \rightarrow_K | t | f | u |
|-----------------|----------|----------|----------|
| t | t | f | u |
| f | t | t | t |
| u | t | u | u |

The abstract algorithm is expressed in terms of a function, \models , which given a bisimulation formula E and a symmetrically and transitively closed relation R , returns **t** only if $R \models E$, and **f** only if $R \not\models E$. From now on, R is assumed

$$\begin{array}{l}
R \models E \text{ is case } E \text{ of} \\
\quad \top \quad : \mathbf{t} \\
\quad \mathbf{F} \quad : \mathbf{f} \\
\quad x \equiv y \quad : \text{if } x R^0 y \text{ then } \mathbf{t} \text{ else } \mathbf{u} \\
E' \wedge E'' : R \models E' \wedge_{\kappa} R \models E'' \\
E' \vee E'' : R \models E' \vee_{\kappa} R \models E'' \\
E' \rightarrow E'' : R \models E' \rightarrow_{\kappa} R' \models E'', \\
\quad \text{where } R' = \begin{cases} (R \cup \{(x, y)\})^{\oplus}, & \text{if } E' \text{ is } x \equiv y \\ R, & \text{if } E' \text{ is } \top \text{ or } \mathbf{F} \end{cases} \\
\forall x: E' : R \setminus x \models E'
\end{array}$$

FIGURE 6. The abstract algorithm.

to be symmetrically and transitively closed. Writing \models infix, the definition is given in figure 6.

Notice that \models is well-defined because we take the symmetric and transitive closure of $(R \cup \{(x, y)\})$.

Given a formula E , the initial call to this function will be $\emptyset \models E$, where \emptyset is the empty relation.

Proving correctness is a matter of proving soundness of \models relative to \models . First, we need a small result linking R_{\equiv} to universal quantification.

Lemma 4 If $R_{\equiv} \setminus x \models E$, then $R_{\equiv} \models \forall x: E$.

Proof Assume that $R_{\equiv} \setminus x \models E$ and let an environment ρ satisfying R_{\equiv} be given. Because $R_{\equiv} \setminus x \subseteq R_{\equiv}$, ρ must satisfy $R_{\equiv} \setminus x$ as well. Now x does not occur in any formula of $R_{\equiv} \setminus x$ so all environments differing from ρ only on the value of x , will then also satisfy $R_{\equiv} \setminus x$. By the assumption each such environment also satisfies E wherefore the original environment ρ satisfies $\forall x: E$. \square

Writing $\models E$ for $\emptyset \models E$, we can now state the correctness of the algorithm.

Theorem 5 (Correctness)

- a) If $\models E = \mathbf{t}$, then $\models E$.
- b) If $\models E = \mathbf{f}$, then $\not\models E$.

Proof Part a) of the theorem follows from the stronger statement

$$\text{if } R \models E = \mathbf{t} \text{ then, } R_{\equiv} \models E$$

which we prove by induction on the structure of E . Assume $R \models E = \mathbf{t}$. We consider the forms of E :

\top, \mathbf{F} : In general, $\Gamma \models \top$, so also $R_{\equiv} \models \top$. The case of \mathbf{F} is trivial, since $R \models \mathbf{F} \neq \mathbf{t}$.

$x \equiv y$: By definition of \equiv , it follows that $R \models x \equiv y = \mathbf{t}$ iff either $x R y$ or $x = y$.

Now, $x R y$ is equivalent to $x \equiv y \in R_{\equiv}$. By (Rep), we get $R_{\equiv} \models x \equiv y$. In the case $x = y$, the situation is really that E is $x \equiv x$. By (Ext) and a) of proposition 3, we directly obtain $R_{\equiv} \models x \equiv x$.

$E' \wedge E''$: By definition, $R \models E' \wedge E'' = \mathbf{t}$ implies $R \models E' = \mathbf{t}$ and $R \models E'' = \mathbf{t}$. By induction, we obtain that $R_{\equiv} \models E'$ and $R_{\equiv} \models E''$. Using (Conj), we obtain that $R_{\equiv} \models E' \wedge E''$.

$E' \vee E''$: Similar, using proposition 2 instead of (Conj).

$E' \rightarrow E''$: By the definition of \rightarrow_K , we must have $R \models E' = \mathbf{f}$ or $R' \models E'' = \mathbf{t}$. The forms of E' :

\top : Then $R' = R$ and because $R \models \top = \mathbf{t}$, it follows that $R \models E'' = \mathbf{t}$. As above we deduce $R_{\equiv} \models E''$. Using (Ext), we get $R_{\equiv}, \top \models E''$ and therefore $R_{\equiv} \models \top \rightarrow E''$ follows by (Imp).

\mathbf{f} : In general $\Gamma, F \models E$, so in particular $R_{\equiv}, F \models E''$. By (Impl), $R_{\equiv} \models F \rightarrow E''$.

$x \equiv y$: We have $R \models x \equiv y = \mathbf{t}$ or $R \models x \equiv y = \mathbf{u}$. In either case, we must have $R' \models E'' = \mathbf{t}$, where $R' = (R \cup \{(x, y)\})^{\oplus}$. By induction, we get $R'_{\equiv} \models E''$, which is the same as $(R_{\equiv}, x \equiv y)^{\oplus} \models E''$. Now any $z \equiv w$ in $(R_{\equiv}, x \equiv y)^{\oplus}$ can be deduced from $R_{\equiv}, x \equiv y$ so by repeated use of (Cut), each of these $z \equiv w$ can be removed from the hypothesis and we finally get $R_{\equiv}, x \equiv y \models E''$. Thus, by (Impl), $R_{\equiv} \models x \equiv y \rightarrow E''$.

$\forall x: E'$: By the induction hypothesis we get that $R_{\equiv} \setminus x \models E'$. The result follows from lemma 4.

Part b) follows similarly from the stronger statement that if $R \models E = \mathbf{f}$, then $R_{\equiv} \models \neg E$. \square

6 The Concrete Algorithm

In this section, we discuss the implementation of the abstract algorithm outlined in section 5. The function \models , which is defined there, closely follows the structure of a formula E . The concrete implementation in this section will follow this structure in exactly the same way. So, the primary task is to find a representation of the relation R such that operations on this relation (union, closure, checking for equivalence, etc.) can be performed efficiently.

The primary operations are to make two variables equivalent and to test whether two variables are already equivalent. This is an instance of the so-called disjoint set problem, which is usually solved using rooted trees [6]. To obtain the best possible performance, path compression (McIlroy and Morris) and union by rank [17] (or similar schemes) are normally used to

obtain an amortized complexity of $O(A^{-1}(n))$ per find operation [17, 19], where A^{-1} is the inverse of the (unary) Ackermann function [1].

However, when processing formulas like $(x \equiv y \rightarrow E) \wedge E'$, we need to first form the union of the equivalence classes of x and y , then process the expression E , and then deunion (undo) the last union before processing E' . Path compressions are impossible to undo without ruining the complexity, so we only use union by rank, and obtain a complexity of $O(\log n)$ per find [18]. In order to undo the unions, each union operation is registered on a stack. In this way, deunions can be done in constant time (unions are still constant time). These three operations, find, union, and deunion, can also be implemented such that the amortized complexity for the find operation becomes $O(\log n / (\log \log n))$. That proposal is from [12]. See [20] for the analysis. However, the size of the overhead is so large that for formulas that we consider (up to approximately 5 Mbytes), this method is slower. For further details on disjoint set implementations, see [13]. We call the structure we use a *union-find-deunion (UFD) structure*.

For formulas without universal quantification, this would be all we would need. However, formulas like $(\forall x: E) \wedge E'$ require that the variable x is freed from previous unions while processing E . Afterwards, for the processing of E' , all the old information on x must be restored. Having to keep track of several versions of variables means that the variables cannot be used directly in the *UFD* structure. Instead, we do the following: at any point during the processing of a formula, each variable, x , has an associated stack of pointers corresponding to the number of active versions of x . In greater detail, when a quantifier construction $\forall x:$ is encountered, a pointer is pushed onto x 's stack. The pointer points to a new item in the *UFD* structure not related to anything, which was previously there. In this way, the old environment can be restored by simply popping the stack.

In order to access the stacks associated with variable names as fast as possible, variable names (along with the pointer to the stacks) are organized in a red-black tree [2, 8], which is one of the efficient implementations of dictionaries with a complexity of $O(\log n)$ per operation, where n is the number of elements in the tree.

To summarize, we use a red-black tree that has variable names as keys and stacks of pointers as values. All these pointers point into a common *UFD* structure. In addition, the *UFD* structure has its own stack of undo information. We refer to the structure consisting of all these other data structures as the *combined* structure. See figure 7.

In the following, we list the operations that the three data structures are assumed to be equipped with. The description is brief as all this is quite well known. However, it seems useful to introduce the names of the operations on the different structures.

A stack is a collection of values, which can be removed from the structure only in the reverse order of which they were inserted. Assume that S is a stack and v is a value. The following operations are supported: $Push(S, v)$,




FIGURE 7. The combined structure.

$Pop(S)$, $Top(S)$, $Empty(S)$, and $InitStack()$.

A dictionary implements a set of pairs (k, v) , where k is a key value from a totally ordered domain and v is any value. We assume that each key value appears at most once in the dictionary. If T is a dictionary, then the following operations are supported: $Insert(T, k, v)$, $Delete(T, k)$, $Member(T, k)$, $LookUp(T, k)$, and $InitTree()$.

A *UFD* structure is a collection of elements some of which may be equivalent with other elements. The following operations are supported: $Union(U, p, q)$, $Find(U, p)$, $Deunion(U)$, and $InitUFD()$. Obviously, the implementation is basically the well-known union-find structure using a stack to save information about the unions.

A Kleene boolean is an implementation of Kleenes three-valued logic. The three Kleene truth-values *TRUE*, *FALSE*, and *UNKNOWN* correspond to **t**, **f**, and **u**, respectively. The operations *Kand*, *Kor*, and *Kimp* implement the operations \wedge_K , \vee_K , and \rightarrow_K as described in the tables of section 5. Furthermore, *Ktu* turns an ordinary boolean into the Kleene boolean *TRUE* if it is true and otherwise into *UNKNOWN*.

The Algorithm

In this section, we present the concrete algorithm, which implements the abstract algorithm from section 5. Basically, this is all about representing the relation R using advanced data structures. We assume that the formula E has a representation in the form of a syntax tree. There are well-developed standard techniques to define and manipulate syntax trees. For clarity, we leave out these details.

Also, to present the crucial parts of the algorithm as clearly as possible, we treat $E_1 \wedge E_2$ and $E_1 \vee E_2$ independently. In reality, as we want to process the formula using pipelining, we should process E_1 first and not until after that has been done can we decide whether a conjunction or a disjunction is been processed. Another reasonable assumption would be to require that the program generating the formula does this using a prefix

```

function check( $E$ : formula)  $\rightarrow$  Kleene boolean;
var
  r: Kleene boolean;
  p,q: pointers; (* into the UFD structure *)
case  $E$  of
  T:           r := TRUE;
  F:           r := FALSE;
   $x \equiv y$ :   r := Ktu(Find(U,Top(LookUp(T,x))) =
                Find(U,Top(LookUp(T,y))));
   $E_1 \wedge E_2$ : r := Kand(check( $E_1$ ),check( $E_2$ ));
   $E_1 \vee E_2$ :  r := Kor(check( $E_1$ ), check( $E_2$ ));
   $T \rightarrow E_1$ : r := check( $E_1$ );
   $F \rightarrow E_1$ : r := TRUE;
   $(x \equiv y) \rightarrow E_1$ :
    p := Find(U,Top(LookUp(T,x)));
    q := Find(U,Top(LookUp(T,y)));
    if p  $\neq$  q then Union(U,p,q);
    r := Kimp(Ktu(p  $\neq$  q),check( $E_1$ ));
    if p  $\neq$  q then Deunion(U);
   $\forall x: E_1$ :   if  $\neg$ Member(T,x) then Insert(T,x,InitStack());
                new(p); Push(LookUp(T,x),p);
                r := check( $E_1$ );
                Pop(LookUp(T,x)); free(p);
                if Empty(LookUp(T,x)) then Delete(T,x);
end;
return r;
end;

```

FIGURE 8. The concrete algorithm.

notation like $\wedge(E_1, E_2)$ and $\vee(E_1, E_2)$.

For simplicity, we assume that the formulas are closed, i.e., they do not have any free variables. This is no serious simplification since free variables can be treated as if they were bound at the outermost level.

The concrete implementation, figure 8, follows the structure of the formula in the same way as \models , except that the call for the left-hand operand of implication is unfolded and incorporated directly into the case analysis.

Before use, T is declared as a red-black tree and properly initialized using `InitTree()`. Similarly, U is declared as a *UFD* structure and initialized by a call to `InitUFD()`.

Correctness

Proposition 6 The combined structure immediately after a call to the function `check` is exactly as it were immediately before the call to `check`.

Proof By induction in the number of calls to the function `check`. The base case is when this number is one, which means that `check` is not called recursively. Thus, we must be in one of the cases \top , F , or $x \equiv y$. The result follows since the combined structure is not altered in any of these cases.

For the induction step, the result follows trivially from the induction hypothesis in the case where E is $E' \wedge E''$, $E' \vee E''$, $\top \rightarrow E'$, or $\text{F} \rightarrow E'$, since the combined structure is not changed.

Assume that E is $(x \equiv y) \rightarrow E'$. By the induction hypothesis, the call `check(E')` leaves the structure unchanged. The claim follows as `Deunion(U)` will undo the last union not yet undone. This must be `Union(U,p,q)`, as the combined structure after the call to `check(E')` is exactly as it were before the call.

Assume that E is $\forall x: E'$. By the induction hypothesis, the call `check(E')` leaves the structure unchanged. Since `LookUp(T,x)` is a stack, the statement `Pop(LookUp(T,x))` will undo the effect of `Push(LookUp(T,x),p)`. Furthermore, if the stack `LookUp(T,x)` is empty, then this stack must have been inserted into T by this current invocation of `check`, so the empty stack should be deleted. \square

Proposition 7 Let F be a bisimulation formula, and let E be a subexpression of F with x and y bound in the context of E . Immediately before the call `check(E)`, the combined structure is an exact representation of R in the corresponding call $R \models E$, i.e.,

$$x R^0 y \Leftrightarrow \text{Find}(U, \text{Top}(\text{LookUp}(T,x))) = \text{Find}(U, \text{Top}(\text{LookUp}(T,y)))$$

Proof By induction in the structure of E . The base case is when $E = F$, in which case both R (and thus also R^0) and the combined structure are empty. For the induction step, we consider all possible forms that E could have.

If E is \top , F , $x \equiv y$, $E' \wedge E''$, $E' \vee E''$, $\top \rightarrow E'$, or $\text{F} \rightarrow E'$, then the combined structure remains unchanged and the same R is used in the recursive application of \models .

Assume that E is $(x \equiv y) \rightarrow E'$. Then \models is called with the relation formed by adding (x, y) to R and taking the symmetric and transitive closure. In the combined structure, if x and y do not already belong to the same equivalence class, then the equivalence classes of x and y are joined. Notice that given the representation of the combined structure and the way it is used, it is automatically closed reflexively, symmetrically, and transitively.

Assume that E is $\forall x: E'$. Then \models is called with the relation formed from R by deleting all pairs that include x . In the combined structure, a new pointer into the *UFD* structure is created and placed on the top of x 's variable stack, thus effectively hiding any pairs involving x ; except that the pair (x, x) will belong to the structure ensuring reflexivity. \square

Lemma 8 The function, `check`, correctly implements \models .

Proof From proposition 6, it follows that a function semantically equivalent to the function `check` can be written by letting the combined structure be a value-passing parameter to `check`. As the two algorithms are structurally equivalent modulo unfolding, it is sufficient to consider the use of the combined structure and R . From proposition 7, it follows that the combined structure is an exact representation of R^0 . \square

Theorem 9 Let E be a bisimulation formula, and let n be the size of E . Then the time-complexity of `check`(E) is $O(n \log n)$.

Proof The algorithm is recursive in the structure of E , and clearly, there are a constant number of statements per symbol in E . These statements either perform constant-time operations, or they operate on one of the data structures. As these are initially empty, and as they share the property that n operations are carried out in time $O(n \log n)$, the result follows. \square

7 Extensions of the Algorithm

In this section, we consider various extensions of the algorithm. For each extension, we sketch the modifications of the abstract algorithm from section 5, and we discuss the correctness issues briefly.

Many more extension than the ones presented here are possible. However, we have decided only to present extensions according to the criteria:

- a) the asymptotic complexity should not change.
- b) the increase in the actual complexity should be very low (less than a factor of 10).
- c) it should still be a one pass algorithm.

It is not hard to deal with constants and through the obvious transformation suggested by the equivalence

$$(E \wedge E') \rightarrow E'' \models E \rightarrow (E' \rightarrow E''),$$

the algorithm can easily cover implication subformulas with conjunctions of predicates to the left. It is straight forward to cope with multiple equivalence relations by letting the function \models work with multiple relations over variables.

The function \models is only able to return **t** (**f**) if the formula is valid (unsatisfiable). However, the algorithm, \models_c , obtained from \models by returning the result of $(R \cup \{(x, y)\})^\oplus \models_c E'$ in case of formulas of the form $(x \equiv y) \rightarrow E'$, is able to deal with *contingent* formulas as well.

Theorem 10 a) If $\models_c E = \mathbf{t}$ then $\models E$. b) If $\models_c E = \mathbf{f}$ then $\not\models E$.

Proof The proof of a) is almost exactly as the corresponding proof of theorem 5. Part b) is proved by showing that if $R \models_c E = \mathbf{f}$ then E it is not satisfied by environments identifying all variables. \square

\models_c is clearly as good as \models . It is also strictly better because $\models x \equiv y \rightarrow \mathbf{F} = \mathbf{u}$ and $\models_c x \equiv y \rightarrow \mathbf{F} = \mathbf{f}$.

8 Simplifications

In this section, we discuss changes to the algorithm with the purpose of outputting a simplified formula equivalent to the original formula. The algorithm should contain the validity checking algorithm as a special case, i.e., if the validity checking algorithm deems a formula valid, then this new algorithm should simplify the formula to \mathbf{T} . Also, we would like the algorithm to fulfill the criteria of the previous section.

Like the \models function, the new function, \models_r , takes as arguments a relation, R , over variables and a first order formula, but now it returns a first order formula instead of a truth-value of three-valued logic. We use the same case analysis, but turn the Kleene truth tables into simplification tables, essentially by replacing \mathbf{u} by the argument formula. Compare with the Kleene truth tables in section 5. However, this is not quite sufficient. In the $\forall x: E'$ case of \models , $\forall x$ is eliminated completely. This cannot be done here when E' does not simplify to \mathbf{T} or \mathbf{F} , so a simplification table for \forall is also needed.

| \wedge_r | T | F | E' |
|------------|-----|---|---------------|
| T | T | F | E' |
| F | F | F | F |
| E | E | F | $E \wedge E'$ |

| \rightarrow_r | T | F | E' |
|-----------------|---|----------------------------|--------------------|
| T | T | F | E' |
| F | T | T | T |
| E | T | $E \rightarrow \mathbf{F}$ | $E \rightarrow E'$ |

| \vee_r | T | F | E' |
|----------|---|-----|-------------|
| T | T | T | T |
| F | T | F | E' |
| E | T | E | $E \vee E'$ |

| $\forall_r x:$ | |
|----------------|-----------------|
| T | T |
| F | F |
| E' | $\forall x: E'$ |

We are now ready to define \models_r .

$$\begin{aligned}
 R \models_r E \text{ is case } E \text{ of } & \quad x \equiv y \quad : \text{ if } x R^0 y \text{ then } \mathbf{T} \text{ else } x \equiv y \\
 & \quad E' \wedge E'' \quad : R \models_r E' \wedge_r R \models_r E'' \\
 & \quad E' \vee E'' \quad : R \models_r E' \vee_r R \models_r E'' \\
 & \quad E' \rightarrow E'' \quad : R \models_r E' \rightarrow_r \mathbf{Upd}(R, E') \models_r E'' \\
 & \quad \forall x: E' \quad : \forall_r x: R \setminus x \models_r E' \\
 & \quad E \quad : E
 \end{aligned}$$

The final case deals with \mathbf{T} and \mathbf{F} and for convenience, we have introduced

an explicit update function:

$$\mathbf{Upd}(R, E) = \begin{cases} (R \cup \{(x, y)\})^\oplus, & \text{if } E \text{ is } x \equiv y \\ R, & \text{otherwise} \end{cases}$$

The simplified formula is logically equivalent with the original as stated in the following correctness theorem.

Theorem 11 E if and only if $\models_r E$.

Proof This follows from the statement below which is proved by induction. We omit the details.

$$R_{\equiv} \rightarrow E \text{ iff } R_{\equiv} \rightarrow (R \models_r E) \quad \square$$

The next proposition expresses that \models_r is at least as good as \models .

Proposition 12 If $\models E = \mathbf{t}(\mathbf{f})$, then $\models_r E = \mathbf{T}(\mathbf{F})$.

The next section contains examples of simplifications using this algorithm.

A straightforward improvement of the simplification algorithm can be obtained from the semantic equivalence

$$E' \wedge E'' \models E' \wedge (E' \rightarrow E''). \quad (1.1)$$

Exploiting the simplification of E' , the conjunction case is changed to:

$$\text{let } E'_r = R \models_r E' \text{ in } E'_r \wedge_r \mathbf{Upd}(R, E'_r) \models_r E''.$$

In this way, the algorithm can simplify $(F \vee x \equiv y) \wedge (x \equiv y \rightarrow F)$, for example, to \mathbf{F} .

Along the same lines, the algorithm can be improved further by using:

$$E \vee E' \models (\neg E) \rightarrow E'.$$

Writing $x \not\equiv y$ for the common occurring formula $x \equiv y \rightarrow \mathbf{F}$, we get as a special case:

$$x \not\equiv y \vee E \models x \equiv y \rightarrow E. \quad (1.2)$$

As the algorithm is formulated now, there is a priori nothing that prevents the algorithm from working with more predicates such as $x \not\equiv y$ and $x \leq y$. In fact, the simplification algorithm is still sound since the new predicates are not simplified and do not give rise to updates of R through $\mathbf{Upd}(-, -)$. However, we can use R to simplify the new predicates in some cases, e.g., for $x \not\equiv y$, we can add the case

$$\text{if } x R^0 y \text{ then } \mathbf{F} \text{ else } x \not\equiv y.$$

All actions are internal, so τ has been omitted from the graphs together with the trivial guarding conditions \top . Before proceeding, we invite the reader to try to see which graphs are bisimilar.

Now, applying \models_r to the bisimulation formula mgb_{g_i, g_j} ($\leftrightarrow mgb_{g_j, g_i}$), we get the table of simplified formulas

| i^j | 1 | 2 | 3 | 4 |
|-------|------------------------------------|--|------------------------------------|---|
| 2 | \top | | | |
| 3 | $x \not\equiv y \vee x \equiv y$ | $x \not\equiv y \vee x \equiv y$ | | |
| 4 | $x \equiv y \wedge E$ | $x \equiv y \wedge E$ | E | |
| 5 | $x \equiv y \wedge x \not\equiv y$ | $(x \equiv y \vee x \equiv y) \wedge x \not\equiv y$ | $x \equiv y \wedge x \not\equiv y$ | $x \equiv y \wedge E \wedge x \not\equiv y$ |

where $E = ((x \equiv y \wedge x \not\equiv y) \vee (x \equiv y \wedge x \equiv y))$, and for sake of readability, $x \equiv y \rightarrow \mathbf{F}$ is written $x \not\equiv y$.

If, instead, we apply \models_r with the modifications corresponding to (1.1), many of the formulas are simplified considerably, some even completely as shown in the following table.

| i^j | 1 | 2 | 3 | 4 |
|-------|----------------------------------|--|--------------|--------------|
| 2 | \top | | | |
| 3 | $x \not\equiv y \vee x \equiv y$ | $x \not\equiv y \vee x \equiv y$ | | |
| 4 | $x \equiv y$ | $x \equiv y$ | $x \equiv y$ | |
| 5 | \mathbf{F} | $(x \equiv y \vee x \equiv y) \wedge x \not\equiv y$ | \mathbf{F} | \mathbf{F} |

If we are interested in knowing whether the graphs are bisimilar under all instantiations of x and y , we can check validity of the universal closure of the formulas, i.e., simplify the universally closed formulas to \top or \mathbf{F} if possible. The result of applying \models_{re} (with the modifications mentioned above) yields the table below. We have omitted the quantifiers in the formulas different from \top and \mathbf{F} .

| i^j | 1 | 2 | 3 | 4 |
|-------|----------------------------------|--|--------------|--------------|
| 2 | \top | | | |
| 3 | $x \not\equiv y \vee x \equiv y$ | $x \not\equiv y \vee x \equiv y$ | | |
| 4 | \mathbf{F} | \mathbf{F} | \mathbf{F} | |
| 5 | \mathbf{F} | $(x \equiv y \vee x \equiv y) \wedge x \not\equiv y$ | \mathbf{F} | \mathbf{F} |

Adding to \models_{re} an extra case for $\not\equiv$, the formula in entry $(i, j) = (5, 2)$ would also simplify to \mathbf{F} , and if the modification suggested from (1.2), i.e., transforming $x \not\equiv y \vee E$ to $x \equiv y \rightarrow E$, is incorporated into the algorithms as well, then the last two entries would also simplify completely, but this time to \top .

Turning to the quantitative aspects of the concrete algorithms, we consider processes defined for $i \geq 0$ by

$$p_{i+2} \xrightarrow{\top, a?x_0} q_{i+1}, \quad q_{i+1} \xrightarrow{\top, a?x_1} r_1^i, \quad r_k^{i+1} \xrightarrow{c_k, a?x_{k+1}} r_{k+1}^i \quad \text{and} \quad r_k^0 \xrightarrow{c_k, b!x_k} \mathbf{0},$$

where c_k is the equality $x_{k-1} = x_k$. Initially two values are unconditionally received on a and then, iteratively, values are received on a provided the two most recently received values are equal. Finally, after i iterations and under the same proviso, the last value is send on b . Similarly, we define primed versions which only differ in that c'_k is the equality $x'_0 = x'_k$. That is, the last value received on a is compared with very first.

In order to give the reader examples of how concrete bisimulation formulas look like, we now describe the most general boolean, $mgb_{r_k^i, r_k'^i}$, characterizing those instantiations (environments) of r_k^i and $r_k'^i$ for which they are bisimilar.

$$mgb_{r_k^{i+1}, r_k'^{i+1}} = \wedge \begin{array}{l} c_k \rightarrow (c'_k \wedge \forall x_{k+1} : \forall x'_{k+1} : x_{k+1} = x'_{k+1} \rightarrow mgb_{r_{k+1}^i, r_{k+1}'^i}) \\ c'_k \rightarrow (c_k \wedge \forall x'_{k+1} : \forall x_{k+1} : x'_{k+1} = x_{k+1} \rightarrow mgb_{r_{k+1}^i, r_{k+1}'^i}) \end{array}$$

$$mgb_{r_k^0, r_k'^0} = \wedge \begin{array}{l} c_k \rightarrow (c'_k \wedge x_k = x'_k \wedge \top) \\ c'_k \rightarrow (c_k \wedge x'_k = x_k \wedge \top) \end{array}$$

The most general booleans for the p 's and q 's are similar to the first formula above, except that the conditions here are \top .

For $2 \leq i \leq 13$, we have measured the average time of five runs of a C implementation of \models_c (\models_r) processing $mgb_{p_{i+2}, p'_{i+2}}$ on a SPARC station ELC. To give a few examples, $mgb_{p_{13+2}, p'_{13+2}}$ of size 5.754 Mb was simplified to **t** (\top) in 77284 (78912) milliseconds. Similarly, $mgb_{p_{13+2}, p'_{13+3}}$, which is 4.477 Mb large, was simplified to **f** in 62958 milliseconds. On average, \models_c and \models_r process input at a rate of about 75 Kbytes per second.

10 REFERENCES

- [1] W. Ackermann. Zum Hilbertschen Aufbau der reellen Zahlen. *Math. Annalen*, 99:118–133, 1928.
- [2] R. Bayer. Symmetric Binary B-Trees. *Acta Inform.*, 1:290–306, 1972.
- [3] S.A. Cook. The Complexity of Theorem-Proving Procedures. In *ACM STOC*, pages 151–158, 1971.
- [4] U.H. Engberg. Simple Symbolic Bisimulations. In preparation.
- [5] U.H. Engberg and K.S. Larsen. Efficient Reduction of Bisimulation Formulas. Preprint 47, Dept. of Math. and Computer Science, Odense University, 1993.
- [6] B.A. Galler and M.J. Fischer. An improved equivalence algorithm. *Comm. ACM*, 7:301–303, 1964.
- [7] M.R. Garey and D.S. Johnson. *Computers and Intractability*. W. H. Freeman, 1979.

- [8] L.J.Guibas and R.Sedgewick. A Dichromatic Framework for Balanced Trees. *IEEE FOCS*, 8–21, 1978.
- [9] M. Hennessy and H. Lin. Symbolic Bismulations. Tech. Rep. 1/92, University of Sussex, 1992. Appear in *Theoretical Computer Science* 138 (1995) 353–389.
- [10] M. Hennessy and H. Lin. Proof systems for message-passing process algebras. *CONCUR '93*, pages 202–216, August 1993.
- [11] G.J. Klir and T.A. Folger. *Fuzzy Sets, Uncertainty, and Information*. Prentice-Hall, 1988.
- [12] H. Mannila and E. Ukkonen. The set union problem with backtracking. *LNCS 226*, 236–243, 1986.
- [13] K. Mehlhorn and A. Tsakalidis. Data Structures. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, chapter 6, pages 301–341. Elsevier Science Publishers, 1990.
- [14] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [15] Z. Schreiber. Verification of Value-Passing Systems. In *First North American Process Algebra Workshop*, pages 9.1–9.20. Tech. Rep. 92-15, Johns Hopkins University, 1992.
- [16] D. Scott. Notes on the formalization of logic. Technical report, Sub-faculty of Phil., Oxford, 1981.
- [17] R.E. Tarjan. Efficiency of a good but not linear set union algorithm. *JACM*, 22:215–225, 1975.
- [18] R.E. Tarjan. *Data Structures and Network Algorithms*. Soc. for Industrial and Applied Math., 1983.
- [19] R.E. Tarjan and J.v.Leeuwen. Worst-Case Analysis of Set Union Algorithms. *JACM*, 31:245–281, 1984.
- [20] J. Westbrook and R.E. Tarjan. Amortized analysis of algorithms for set union with backtracking. *SIAM J. Comput.*, 18(1):1–11, 1989.