Real-Time Systems, , 1–21 ()
© Kluwer Academic Publishers, Boston. Manufactured in The Netherlands.

Compact Data Structures and State-Space Reduction for Model-Checking Real-Time Systems

KIM G. LARSEN kgl@cs.auc.dk
BRICS, Aalborg University, Denmark.

FREDRIK LARSSON, fredrikl@docs.uu.se
Department of Information Technology, Uppsala University, Sweden.

PAUL PETTERSSON paupet@docs.uu.se
Department of Information Technology, Uppsala University, Sweden.

WANG YI yi@docs.uu.se
Department of Information Technology, Uppsala University, Sweden.

Abstract. During the past few years, a number of verification tools have been developed for real-time systems in the framework of timed automata. One of the major problems in applying these tools to industrial-sized systems is the huge memory-usage for the exploration of the state-space of a network (or product) of timed automata, as the model-checkers must keep information about not only the control structure of the automata but also the clock values specified by clock constraints.

In this paper, we present a compact data structure for representing clock constraints. The data structure is based on an $\mathcal{O}(n^3)$ algorithm which, given a constraint system over real-valued variables consisting of bounds on differences, constructs an equivalent system with a *minimal* number of constraints. In addition, we have developed an on-the-fly reduction technique to minimize the space-usage. Based on static analysis of the control structure of a network of timed automata, we are able to compute a set of symbolic states that cover all the dynamic loops of the network in an on-the-fly searching algorithm, and thus ensure termination in reachability analysis.

The two techniques and their combination have been implemented in the tool UPPAAL. Our experimental results demonstrate that the techniques result in truly significant space-reductions: for six examples from the literature, the space saving is between 75% and 94%, and in (nearly) all examples time-performance is improved. Noteworthy is also the observation that the two techniques are completely orthogonal.

Keywords: Real-Time Systems, Model Checking, Design Tool, Formal Specification and Verification, Timed Automata

1. Introduction

Reachability analysis has been one of the most successful methods for automated analysis of concurrent systems. Many verification problems e.g. invariant checking can be solved by means of reachability analysis. It can in many cases also be used for checking whether a system described as an automaton satisfies a requirement specification formulated e.g. in linear temporal logic, by converting the requirement to an automaton and thereafter checking whether the parallel composition of the system and requirement automata can reach certain annotated states [32, 21, 2]. However, the major problem in applying reachability analysis is the potential combinatorial

```
PASSED:= {}
WAITING:= {(l_0, D_0)}
repeat
    begin
        get (l, D) from WAITING
        if (l, D) ⊨ φ then return "YES"
        else if D ⊈ D' for all (l, D') ∈ PASSED then
            begin
                add (l, D) to PASSED
                SUCC:={(l_s, D_s) | (l, D) ⇝ (l_s, D_s) and D_s ≠ ∅}
                for all (l_s, D_s) in SUCC do
                    put (l_s, D_s) to WAITING
            end
    end
until WAITING={}
return "NO"
```

Figure 1. An Algorithm for Symbolic Reachability Analysis.

explosion of state spaces. To attack this problem, various symbolic and reduction techniques have been put forward over the last decade to efficiently represent state space and to avoid exhaustive state space exploration (e.g. [11, 17, 31, 12, 13, 16, 5]); such techniques have played a crucial role for the successful development of verification tools for finite-state systems.

In the last few years, new verification tools have been developed, for the class of infinite-state systems known as timed systems [19, 14, 9]. Notably the verification engines of most tools in this category are based on reachability analysis of timed automata following the pioneering work of Alur and Dill [4]. A timed automaton is an extension of a finite automaton with a finite set of real-valued clock-variables. The foundation for decidability of reachability problems for timed automata is Alur and Dill's region technique, by which the infinite state space of a timed automaton due to the density of time, may effectively be partitioned into finitely many equivalence classes i.e. *regions* in such a way that states within each class will always evolve to states within the same classes. However, reachability analysis based on the region technique is practically infeasible due to the potential state explosions arising from not only the control-structure (as for finite-state systems) but also the region space [23].

Efficient data structures and algorithms have been sought to represent and manipulate timing constraints over clock variables (e.g. by Difference Bounded Matrices [7, 15], or Binary Decision Diagrams [11, 6]) and to avoid exhaustive state space

exploration (e.g. by application of partial order reductions [17, 31, 26] or compositional methods [5, 23]). One of the main achievements in these studies is the symbolic technique [15, 33, 20, 34, 23], that converts the reachability problem to that of solving simple constraints. The technique can be simply formulated in an abstract reachability algorithm[1] as shown in Figure 1. The algorithm is to check whether a timed automaton may reach a state satisfying a given state formula φ. It explores the state space of the automaton in terms of *symbolic states* in the form (l, D) where l is a control-node and D is a constraint over clocks variables.

We observe that several operations of the algorithm are critical for efficient implementations. Firstly, the algorithm depends heavily on the test operations for checking the inclusion $D \subseteq D'$ (i.e. the inclusion between the solution sets of D, D') and the emptiness of D_s in constructing the successor set SUCC of (l, D). Clearly, it is important to design efficient data structures and algorithms for the representation and manipulation of clock constraints. One such well-known data structure is that of DBM (*Difference Bounded Matrix*), which offers a canonical representation for constraint systems. It has been successfully used in several real-time verification tools, e.g. UPPAAL [9] and KRONOS [14]. A DBM representation is in fact a weighted directed graph where the vertices correspond to clocks (including a zero-clock) and the weights on the edges stand for the bounds on the differences between pairs of clocks [7, 15, 33]. As it gives an explicit bound for the difference between each pair of clocks, its space-usage is in the order of $\mathcal{O}(n^2)$ where n is the number of clocks. However, in practice it often turns out that most of these bounds are redundant.

In this paper, we present a compact data structure for DBM, which provides *minimal* and *canonical* representations of clock constraints and also allows for efficient inclusion checks. We have developed an $\mathcal{O}(n^3)$ algorithm that given a DBM constructs a minimal number of constraints equivalent to the original constraints represented by the DBM (i.e. with the same solution set). The algorithm is essentially a minimization algorithm for weighted directed graphs, and hence solves a problem of independent interest. Note that the main global data structure of the algorithm in Figure 1 is the passed list (i.e. PASSED) holding the explored states. In many cases, it will store all the reachable symbolic states of the automaton. Thus, it is desirable that when saving a (symbolic) state in the passed list, we save the (often substantially smaller) minimal constraint system. The minimal representation also makes the inclusion-checking of the algorithm more efficient. Our experimental results demonstrate truly significant space-savings as well as better time-performance (see statistics in section 5).

In addition to the *local* reduction technique above, which is to minimize the space-usage of each individual symbolic state, as the second contribution of this paper, we have developed a *global* reduction technique to reduce the total number of states to save in the global data structure, i.e. the passed list. It is completely orthogonal to the local technique. In the abstract algorithm of Figure 1, we notice the step of saving the new encountered state (l, D) in the passed list when the inclusion-checking for $D \subseteq D'$ fails (i.e. $D \not\subseteq D'$). Its purpose is first of all to guarantee termination but also to avoid repeated exploration of states that have several predecessors. However, this is not necessary if all the predecessors of (l, D) are already

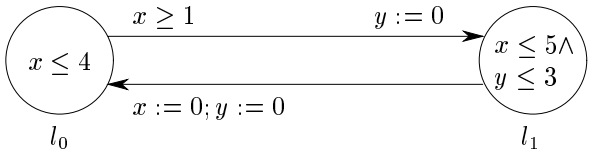

Figure 2. A Timed Automaton.

present in the passed list. In fact, to ensure termination, it suffices to save only one state for each dynamic loop. An improved on-the-fly reachability algorithm according to the global reduction strategy has been implemented in UPPAAL [9] based on static analysis of the control structure of timed automata. Our experimental results demonstrate significant space-savings and also better time-performance (see statistics in section 5).

The outline of this paper is as follows: In the next section we review the semantics of timed automata and the notion of Difference Bounded Matrix (DBM) for clock constraints. Section 3 presents the compact data structure for DBM and the local reduction technique (i.e. the minimization algorithm for weighted directed graphs). Section 4 is devoted to develop the global reduction technique based on control structure analysis. Section 5 presents our experimental results for both techniques and their combination. Section 6 concludes the paper.

2. Preliminaries

2.1. Timed Automata

The model of timed automata was first introduced in [4] and has since then established itself as a standard model for real-time systems. For the reader not familiar with the notion of timed automata we give a short informal description.

Consider the timed automaton of Figure 2. It has two control nodes l_0 and l_1 and two real-valued clocks x and y. A *state* of the automaton is of the form (l, s, t), where l is a control node, and s and t are non-negative reals giving the value of the two clocks x and y. A control node is labelled with a condition (the invariant) on the clock values that must be satisfied for states involving this node. Assuming that the automaton starts to operate in the state $(l_0, 0, 0)$, it may stay in node l_0 as long as the invariant $x \leq 4$ of l_0 is satisfied. During this time the values of the clocks increase synchronously. Thus from the initial state, all states of the form (l_0, t, t), where $t \leq 4$, are reachable. The edges of a timed automaton may be decorated with a condition (guard) on the clock values that must be satisfied in order to be enabled. Thus, only for the states (l_0, t, t), where $1 \leq t \leq 4$, is the edge from l_0 to l_1 enabled. Additionally, edges may be labelled with simple assignments reseting clocks. For example, when following the edge from l_0 to l_1 the clock y is reset to 0 leading to states of the form $(l_1, t, 0)$, where $1 \leq t \leq 4$.

In general, a timed automaton is a standard finite-state automaton extended with a finite collection $\mathcal{C}$ of real-valued clocks ranged over by x, y etc. We use $\mathcal{B}(\mathcal{C})$ ranged over by g (and latter D), to stand for the set of formulas that can be an atomic constraint of the form: $x \sim n$ or $x - y \sim n$ for $x, y \in \mathcal{C}$, $\sim \in \{\leq, \geq\}^2$ and n being a natural number, or a conjunction of such formulas. Elements of $\mathcal{B}(\mathcal{C})$ are called *clock constraints* or *constraint systems* over $\mathcal{C}$.

Definition 1. [Timed Automata] A timed automaton A over clocks $\mathcal{C}$ is a tuple $\langle N, l_0, \longrightarrow, I\rangle$ where N is a finite set of nodes (control-nodes), l_0 is the initial node, $\longrightarrow \subseteq N \times \mathcal{B}(\mathcal{C}) \times 2^{\mathcal{C}} \times N$ corresponds to the set of edges, and finally, $I : N \mapsto \mathcal{B}(\mathcal{C})$ assigns invariants to nodes. In the case, $\langle l, g, r, l'\rangle \in \longrightarrow$, we write $l \xrightarrow{g,r} l'$. □

Formally, we represent the values of clocks as functions (called clock assignments) from $\mathcal{C}$ to the non-negative reals $\mathbb{R}_+$. We denote by $\mathbb{R}_+^{\mathcal{C}}$ the set of clock assignments for $\mathcal{C}$. A semantical *state* of an automaton A is now a pair (l, u), where l is a node of A and u is a clock assignment for $\mathcal{C}$, and the semantics of A is given by a transition system with the following two types of transitions (corresponding to delay-transitions and edge-transitions):

- $(l, u) \longrightarrow (l, u \oplus d)$ if $I(l)(u)$ and $I(l)(u \oplus d)$
- $(l, u) \longrightarrow (l', u')$ if there exist g and r such that $l \xrightarrow{g,r} l'$, $g(u)$ and $u' = r[u]$

where for $d \in \mathbb{R}_+$, $u \oplus d$ denotes the time assignment which maps each clock x in $\mathcal{C}$ to the value $u(x) + d$, and for $r \subseteq \mathcal{C}$, $r[u]$ denotes the assignment for $\mathcal{C}$ which maps each clock in r to the value 0 and agrees with u over $\mathcal{C} \backslash r$.

Clearly, the semantics of a timed automaton yields an infinite transition system, and is thus not an appropriate basis for decision algorithms. However, efficient algorithms may be obtained using a finite-state *symbolic* semantics based on *symbolic states* of the form (l, D), where $D \in \mathcal{B}(\mathcal{C})$ [20, 34]. We shall consider a clock constraint as a set of clock assignments and use $u \in D$ to stand for u satisfied D.

The symbolic counterpart to the standard semantics is given by the following two (fairly obvious) types of symbolic transitions:

- $(l, D) \rightsquigarrow \Big(l, (D \wedge I(l))^{\uparrow} \wedge I(l)\Big)$
- $(l, D) \rightsquigarrow \Big(l', r(g \wedge D)\Big)$ if $l \xrightarrow{g,r} l'$

where $D^{\uparrow} = \{u \oplus d \mid u \in D \wedge d \in \mathbb{R}_+\}$ and $r(D) = \{r[u] \mid u \in D\}$. It may be shown that $\mathcal{B}(\mathcal{C})$ (the set of clock constraints) is closed under these two operations (and $\wedge$) [15]. Moreover, the symbolic semantics characterize the standard semantics in the sense that, whenever $u \in D$ and $(l, D) \rightsquigarrow (l', D')$ then $(l, u) \longrightarrow (l', u')$ for $u' \in D'$.

Finally, we introduce the notion of networks of timed automata [34, 23]. A network is the parallel composition of a finite set of automata for a given synchronization function. To illustrate the on-the-fly verification technique, we only need to study the case dealing with interleaving, that is, the network of automata $A_1 \dots A_n$, is the Cartesian product of A_i's. Assume a vector l of control nodes. We shall use $l[i]$ to stand for the ith element of l and $l[l'_i/l_i]$ for the vector where the ith element l_i of l is replaced by l'_i. A control node (i.e. *control vector*) l of a network $A_1 \dots A_n$ is a vector where $l[i]$ is a node of A_i and the invariant $I(l)$ of l is the conjunction of $I(l[1]) \dots I(l[n])$. The symbolic semantics of networks is given in terms of control vectors by the following two types of symbolic transitions:

- $(l, D) \rightsquigarrow \Big(l, (D \wedge I(l))^{\uparrow} \wedge I(l)\Big)$
- $(l, D) \rightsquigarrow \Big(l[l'_i/l_i], r(g \wedge D)\Big)$ if $l_i \xrightarrow{g,r} l'_i$

In the later case, we shall say that the symbolic transition is derived by the edge $l_i \xrightarrow{g,r} l'_i$.

2.2. Difference Bounded Matrices & Shortest-Path Closure

To utilize the symbolic semantics of (networks of) timed automata algorithmically, as for example in the reachability algorithm of Figure 1, it is important to design efficient data structures and algorithms for the representation and manipulation of clock constraints.

One such well-known data structure is that of difference bounded matrices (DBM, see [7, 15]), which offers a canonical representation for constraint systems. A DBM representation of a constraint system D is simply a weighted, directed graph, where the vertices correspond to the clocks of C and an additional zero-vertex 0. The graph has an edge from x to y with weight m provided $y - x \leq m$ is a constraint of D. Similarly, there is an edge from 0 to x with weight m, whenever $x \leq m$ is a constraint of D [3]. As an example, consider the constraint system E over $\{x_0, x_1, x_2, x_3\}$ being a conjunction of the atomic constraints $x_0 - x_1 \leq 3$, $x_3 - x_0 \leq 5$, $x_3 - x_1 \leq 2$, $x_2 - x_3 \leq 2$, $x_2 - x_1 \leq 10$, and $x_1 - x_2 \leq -4$. The graph representing E is given in Figure 3 (a).

In general, the same set of clock assignments may be described by several constraint systems (and hence graphs). To test for inclusion between constraint systems D and D' [4], which we recall is essential for the termination of the reachability algorithm of Figure 1, it is advantageous if D is *closed under entailment* in the sense that no constraint of D can be strengthened without reducing the solution set. In particular, for D a closed constraint system, $D \subseteq D'$ holds if and only if for any constraint in D' there is a constraint in D at least as tight; i.e. whenever $(x - y \leq m') \in D'$ then $(x - y \leq m) \in D$ for some $m \leq m'$. Thus, closedness provides a canonical representation, as two closed constraint systems describe the same solution set precisely when they are identical. To close a constraint system D

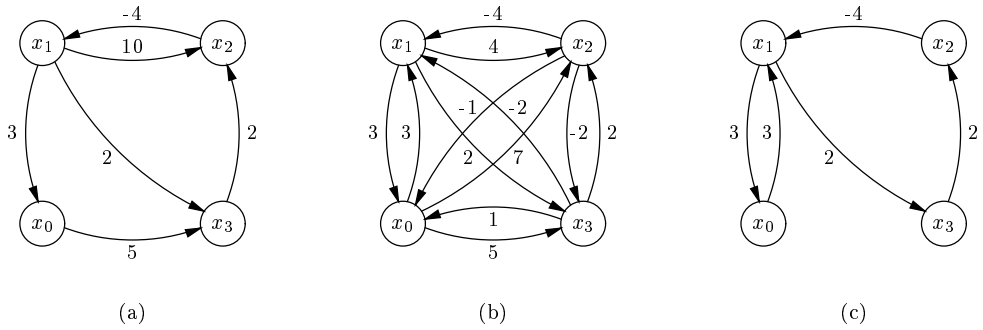

Figure 3. Graph for E (a), its shortest-path closure (b), and shortest-path reduction (c).

amounts to derive the shortest-path closure for its graph and can thus be computed in time $\mathcal{O}(n^3)$, where n is the number of clocks of D. The graph representation of the closure of the constraint system E from Figure 3 (a) is given in Figure 3 (b). The emptiness-check of a constraint system D simply amounts to checking for negative-weight cycles in its graph representation. Finally, given a closed constraint system D the operations $D^{\uparrow}$ and $r(D)$ may be performed in time $\mathcal{O}(n)$.

3. Minimal Constraint Systems & Shortest Path Reductions

For the reasons stated above a matrix representation of constraint systems in closed form is an attractive data structure, which has been successfully employed by a number of real-time verification tools, e.g. UPPAAL [9] and KRONOS [14]. As it gives an explicit (tightest) bound for the difference between each pair of clocks (and each individual clock), its space-usage is of the order $\mathcal{O}(n^2)$. However, in practice it often turns out that most of these bounds are redundant, and the reachability algorithm of Figure 1 is consequently hampered in two ways by this representation. Firstly, the main data structure PASSED, will in many cases store all the reachable symbolic states of the automaton. Thus, it is desirable, that when saving a symbolic state in the PASSED-list, we save a representation of the constraint system with as few constraints as possible. Secondly, a constraint system D added to the PASSED-list is subsequently only used in checking inclusions of the form $D' \subseteq D$. Recalling the method for inclusion-check from the previous section, we note that (given D' is closed) the time-complexity of the inclusion-check is linear in the number of constraints of D. Thus, again it is advantageous for D to have as few constraints as possible.

In the following subsections we shall present an $\mathcal{O}(n^3)$ algorithm, which given a constraint system constructs an equivalent reduced system with the minimal number of constraints. The reduced constraint system is canonical in the sense that two constrain systems with the same solution set give rise to identical reduced sys-

tems. The algorithm is essentially a minimization algorithm for weighted directed graphs. Given a weighted, directed graph with n vertices, it constructs in time $\mathcal{O}(n^3)$ a reduced graph with the minimal number of edges having the same shortest path closure as the original graph. Figure 3 (c) shows the minimal graph of the graphs in Figure 3 (a) and (b), which is computed by the algorithm.

3.1. Reduction of Zero-Cycle Free Graphs

A weighted, directed graph G is a structure (V, E_G), where V is a finite set of vertices and E_G, is a partial function from $V \times V$ to Z (the integers). The domain of E_G constitutes the edges of G, and when defined, $E_G(x, y)$ gives the weight of the edge between x and y. We assume that $E_G(x, x) = 0$ for all vertices x, and that G has no cycles with negative weight[5].

Given a graph G, we denote by G^C the *shortest-path closure* of G, i.e. $E_{G^C}(x, y)$ is the length of the shortest path from x to y in G. A *shortest-path reduction* of a graph G is a graph G^R with the minimal number of edges such that $(G^R)^C = G^C$.

The key to reduce a graph is obviously to remove *redundant edges*, where an edge (x, y) is redundant if there exist an alternative path from x to y whose (accumulated) weight does not exceed the weight of the edge itself. For example, in the graph of Figure 3 (a), the edge (x_1, x_2) is clearly redundant as the accumulated weight of path $(x_1, x_0), (x_0, x_3), (x_3, x_2)$ has a weight (10) not exceeding the weight of the edge itself (also 10). The path $(x_1, x_3), (x_3, x_2)$ makes also the edge (x_1, x_2) redundant. Being redundant, the edge (x_1, x_2) may be removed without changing the shortest-path closure. We shall use $G \setminus (x_1, x_2)$ to denote the result of removing the edge (x_1, x_2) from the graph G.

Now, consider the edge (x_1, x_2) in the graph of Figure 3 (b). Clearly, the edge is redundant as the path $(x_1, x_3), (x_3, x_2)$ has equal weight. Similarly, the edge (x_3, x_2) is redundant as the path $(x_3, x_1), (x_1, x_2)$ has equal weight. However, though redundant, we cannot just remove the two edges (x_1, x_2) and (x_3, x_2) as removal of one clearly requires the presence of the other. In fact, all edges between the vertices x_1, x_2 and x_3 are redundant, but obviously we cannot remove them all simultaneously. The key explanation of this complicating phenomena is that x_1, x_2, x_3 constitutes a cycle with length zero (a *zero-cycle*). However, for zero-cycle free graphs the situation is the simplest possible:

LEMMA 1 *Let G_1 and G_2 be zero-cycle free graphs such that ${G_1}^C = {G_2}^C$. If there is an edge $(x, y) \in G_1$ such that $(x, y) \notin G_2$, then $(G_1 \setminus \{(x, y)\})^C = {G_1}^C = {G_2}^C$.*

Proof: Let α denote the edge (x, y) and let m be the weight of α in G_1. We will show that there is an alternative path in G_1 *not* using α with weight no more than m. From this fact the Lemma obviously follows.

As ${G_1}^C = {G_2}^C$, the shortest path from x to y in G_2 has weight no more than m. As $\alpha \notin G_2$, this path must visit some vertex z different from x and y. Now let m_1 be the shortest path-weight from x to z and let m_2 be the shortest path-weight from z to y; note that G_1 and G_2 agrees on m_1 and m_2, as they have the same shortest-path closure. Then clearly, $m \geq m_1 + m_2$.

Now assume that the shortest path in G_1 from x to z uses $\alpha = (x, y)$. Then, as a sub-path, G_1 will be a path from y to z. Since G_1 also has a path from z to y, it follows that G_1 will have a cycle from y via z back to y. The weight of this cycle can be argued to be no more than $(m_1 - m) + m_2$. However, as $m \geq m_1 + m_2$ and there are no negative cycles, this cycle must have weight 0 contradicting the assumption that G_1 is zero-cycle free.

Similarly, a contradiction with the zero-cycle free assumption of G_1 is obtained, if the shortest path in G_1 from z to y uses α. thus we can conclude that there is an path from x to y not using α with length no greater than m. □

From the above Lemma it follows immediately that all redundant edges of a zero-cycle free graph may be removed without affecting the closure. On the other hand, removal of an edge which is not redundant will of course change the closure of the graph, and must be present in any graph with the same closure. Thus the following theorem follows:

THEOREM 1 *Let G be a zero-cycle free graph, and let $\{\alpha_1, \ldots, \alpha_2\}$ be the set of redundant edges of G. Then $G^R = G^C \setminus \{\alpha_1, \ldots, \alpha_k\}$.*

Proof: Follows from Lemma 1. □

From an algorithmic point of view, redundancy of edges is easily determined given the closure G^C of a graph G as only path of length 2 needs to be considered: An edge (x, y) is redundant precisely when there is a vertex z ($\neq x, y$) such that $E_{G^C}(x, y) \geq E_{G^C}(x, z) + E_{G^C}(z, y)$. Thus for zero-cycle free graphs computing G^R is $\mathcal{O}(n^3)$.

3.2. Reduction of Negative-Cycle Free Graphs

For general graphs (without negative cycles) our reduction construct relies on a partitioning of the vertices according to zero-cycles. We say that two vertices x and y are *equivalent* or *zero-equivalent*, if there is a zero-cycle containing them both. We write $x \equiv y$ in this case. Given the closure G^C of a graph G, it is extremely easy to check for zero-equivalence: $x \equiv y$ holds precisely when $E_{G^C}(x, y) = -E_{G^C}(y, x)$. Thus, in the graphs of Figure 3 (a) and (b), $\equiv$ partitions the vertices into the two classes $\{x_0\}$ and $\{x_1, x_2, x_3\}$.

To obtain a canonical reduction, we assume that the vertices of G are ordered by assigning them indices as $x_1, x_2, \ldots, x_n$. The equivalence $\equiv$ now induces a natural transformation $G_\equiv$ on the graph G:

Definition 2. Given a graph G, the vertices of the graph $G_\equiv$ are $\equiv$-equivalence classes, denoted E_k, of G. There is an edge between the classes E_i and E_j ($i \neq j$) if for some $x \in E_i$ and $y \in E_j$ there is an edge in G between x and y. The weight of this edge is $E_{G^C}(E_i^{min}, E_j^{min})$, where E^{min} is the vertex in E with the smallest index. □

Thus, the distance between E_i and E_j in $G_{\equiv}$ is the weight of the shortest path in G between the elements of E_i and E_j with smallest index. It is obvious that $G_{\equiv}$ is a zero-cycle free graph. It is also easy to see that $G_{1\equiv} = G_{2\equiv}$ if ${G_1}^C = {G_2}^C$. Let H be the graph of Figure 3 (a). Then $H_{\equiv}$ will have vertices $E_0 = \{x_0\}$ and $E_1 = \{x_1, x_2, x_3\}$. The two vertices are connected by two edges both having weight 3.

The following provides a dual to the operator of Definition 2:

Definition 3. Let F be a graph with vertices being $\equiv$-equivalence classes with respect to a graph $G = (V, E_G)$. Then the expansion of F is a graph F^+ with vertices V and with weight satisfying:

- For any multi-member equivalence class[6] $\{z_1 < z_2 < \cdots < z_k\}$ of F, F^+ contains a single cycle $z_1, z_2, \ldots, z_k, z_1$, with the weight of the edge (z_i, z_{i+1}) being the weight of the shortest path from z_i to z_{i+1} in G.
- Whenever (E_i, E_j) is an edge in F with weight m, then F^+ will have an edge from E_i^{min} to E_j^{min} with weight m. □

We are now ready to state the main Theorem giving the shortest-path reduction construct for arbitrary negative-cycle free graphs:

THEOREM 2 *Let G be negative-cycle free graph. Then the shortest-path reduction G^R of G is given by the graph $({G_{\equiv}}^R)^+$, i.e. $G^R = ({G_{\equiv}}^R)^+$.*

Proof: We show: (1) that $({G_{\equiv}}^R)^+$ is a candidate for a shortest-path reduction of G in the sense that $({G_{\equiv}}^R)^+ = G^C$, and (2) that $({G_{\equiv}}^R)^+$ is minimal.

1. We first prove that $({G_{\equiv}}^R)^+ = G^C$. As all edges (x, y) of $({G_{\equiv}}^R)^+$ have weight of the form $E_{G^C}(x, y)$, it follows that for any path in $({G_{\equiv}}^R)^+$ there is a path in G with same weight.

   Now consider an edge (x, y) of G. We will demonstrate that there is a path in $({G_{\equiv}}^R)^+$ with no greater weight.

   - If $x = E_i^{min}$ and $y = E_j^{min}$ for two $\equiv$-classes E_i and E_j, it follows that $E_{G_{\equiv}}(E_i, E_j) \leq E_G(x, y)$. Furthermore, due to the property of reduction construction, there is a path in ${G_{\equiv}}^R$ between E_i and E_j with weight no greater than $E_{G_{\equiv}}(E_i, E_j)$. The same path, but now between the nodes with the minimal indices of the $\equiv$-classes, can be found in $({G_{\equiv}}^R)^+$. Thus, there is a path in $({G_{\equiv}}^R)^+$ with weight no greater than $E_G(x, y)$.
   - If $x, y \in E_i$ for some $\equiv$-class E_i, an easy argument gives that $E_{({G_{\equiv}}^R)^+}(x, y) = E_{G^C}(x, y) \leq E_G(x, y)$.
   - Consider the case when $x \in E_i$ and $y \in E_j$ for two different $\equiv$-classes, and assume that $E_G(x, y) = m$.

Let $m_1 = E_{(G_{\equiv^R})^+}(x, E_i^{min})$, $m_2 = E_{(G_{\equiv^R})^+}(E_i^{min}, E_j^{min})$, and $m_3 = E_{(G_{\equiv^R})^+}(E_j^{min}, y)$. Note that by the reduction construction $m_2 \leq E_{G^C}(E_i^{min}, E_j^{min})$. Then there is a path in $(G_{\equiv}{}^R)^+$ from x to y via E_i^{min} and E_j^{min} with weight $m_1 + m_2 + m_3$. Now, if $m < m_1 + m_2 + m_3$, there is a path in G from E_i^{min} to E_j^{min} of weight $m - m_1 - m_3 < m_2$ contradicting that m_2 is the weight of the shortest path in G between E_i^{min} and E_j^{min}. Thus the path x, E_i^{min}, E_j^{min}, y in $(G_{\equiv}{}^R)^+$ has weight no greater than the edge (x, y) in G.

2. Next we prove that $(G_{\equiv}{}^R)^+$ has minimal number of edges by showing that whenever $H^C = G^C$ then H has at least as many edges as $(G_{\equiv}{}^R)^+$.

   As $H^C = G^C$, H and G induces the same $\equiv$-equivalence relation on the same zero-length cycles. Obviously the fewest edges that will identify k (> 1) vertices, with respect to $\equiv$ is k. Hence, $(G_{\equiv}{}^R)^+$ uses a minimal number of edges between vertices in the same $\equiv$-equivalence class.

   Now let (E_i^{min}, E_j^{min}) be an edge in $(G_{\equiv}{}^R)^+$ with weight m. We claim that H must have at least one edge from E_i to E_j.

   Assume that this is not the case. Then, as $H^C = G^C$, there must be a path in H from E_i^{min} to E_j^{min} as shown in Figure 4 such that $m = \sum_{i=0}^{k+2} v_i + \sum_{i=0}^{k+1} w_i$.

   Now let $m_0 = E_{(G_{\equiv^R})^+}(E_i^{min}, E_0^{min})$, $m_1 = E_{(G_{\equiv^R})^+}(E_0^{min}, E_1^{min})$, $\ldots$, $m_{k+1} = E_{(G_{\equiv^R})^+}(E_k^{min}, E_i^{min})$ (illustrated with dashed lines in Figure 4). Then $m_0 \leq v_0 + w_0 + v_1'$, $m_1 \leq v_1'' + w_1 + v_2'$, $\ldots$, $m_{k+1} \leq v_{k+1}'' + w_{k+1} + v_{k+2}$, where $v_1 = v_1' + v_1''$, $v_2 = v_2' + v_2''$, $\ldots$, $v_{k+1} = v_{k+1}' + v_{k+1}''$.

   It follows that $\sum_{i=0}^{k+1} m_i \leq m$. Hence (E_i, E_j) is redundant in $(G_{\equiv}{}^R)^+$ and can be removed, contradicting Lemma 1. □

First, note that the above construction of $(G_{\equiv}{}^R)^+$ is well-defined as $G_{\equiv}$ is a zero-cycle free graph and the reduction construction of Theorem 1 thus applies. Given the closure G^C of G the constructions of Definitions 2 and 3 can be computed in $\mathcal{O}(n^2)$. Since G^R is computed from G in $\mathcal{O}(n^3)$, it follows that also $(G_{\equiv}{}^R)^+$ can be constructed in $\mathcal{O}(n^3)$. Now applying the above construction to the graph H of Figure 3 (a), we first note that $H_{\equiv}{}^R = H_{\equiv}$ as $H_{\equiv}$ has no redundant edges. Expanding $H_{\equiv}$ with respect to the vertex ordering $x_0 < x_1 < x_2 < x_3$ gives the graph of Figure 3 (c), which according to Theorem 2 above is the shortest-path reduction of H.

Experimental results show that the use of minimal constrain systems (obtained by the above shortest-path reduction algorithm) as a compact data structure leads to truly significant space-savings in practical reachability analysis of timed systems: the space-savings are in the range 68–85%. We refer to Section 5 for more details.

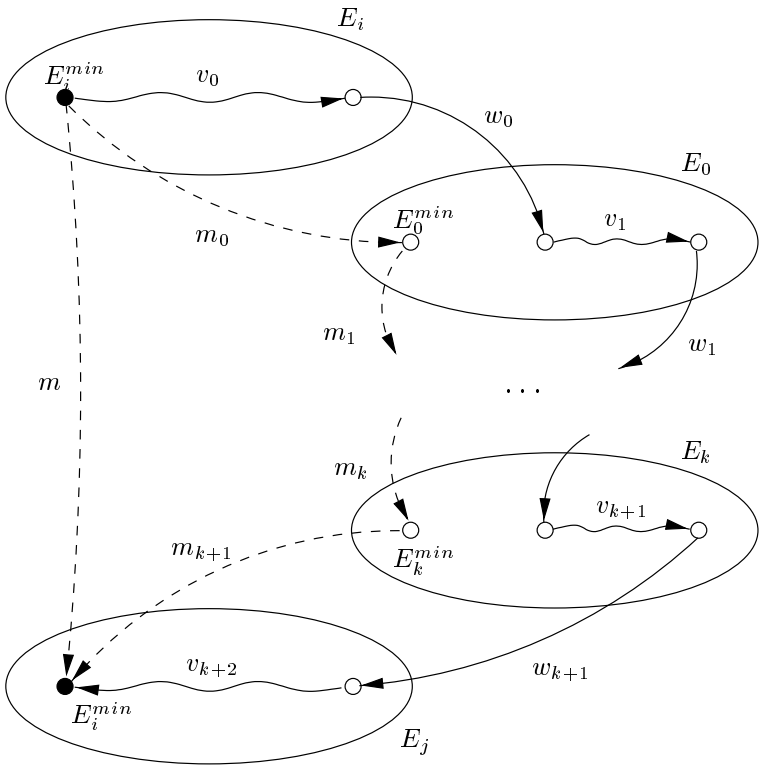

Figure 4. A path in the graph H.

4. Global Reductions and Control Structure Analysis

The preceding section is about *local* reductions in reachability analysis in the sense that the technique developed is for *each* individual symbolic state. In this section, we shall develop a *global* reduction technique to reduce the total number of symbolic states to save in the *global* data structure i.e. the passed list.

4.1. Potential Space-Reductions

We recall the standard reachability analysis algorithm for finite graphs (see e.g. [27]). It is similar to the one in Figure 1, but simpler as no constraints but only control nodes are involved. The algorithm repeats three main operations: *examining* every new encountered node (to see if it is in the passed list), *exploring* the new encountered nodes (computing all their successors for further analysis), and *saving* the explored nodes in the passed list until all reachable nodes are present in the list (i.e. all new encountered nodes are already in the passed list).

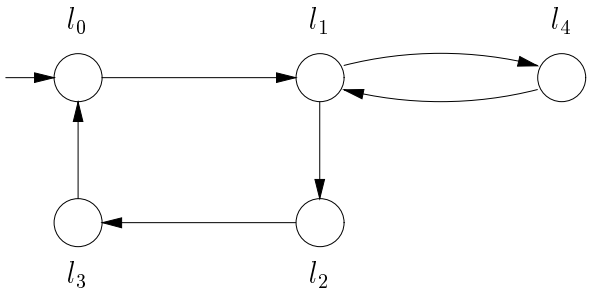

Figure 5. Illustration of Space-Reduction.

Note that the saving of an explored node is to ensure termination and also to avoid repeated exploration of nodes with more than one incoming edge. However it is not necessary to save all reachable nodes. Consider for example, the simple graph in Figure 5 with initial node l_0. Clearly, there is no need to save node l_2, l_3 or l_4 as they will be visited only once if l_1 is present in the passed list.

In fact, to guarantee termination on a finite graph, it is sufficient to save only one node for each cycle in the graph. For example, as l_1 covers the two cycles of the graph in Figure 5, in addition to l_2, l_3, and l_4, it is not necessary to save l_0 either. In general, for a finite graph, there is a minimal number of nodes to save in the passed list in order to guarantee termination. However the trade-off of the space-saving strategy may be increased time-consumption. Consider the same graph of Figure 5. If node l_0 is not present in the passed list, it will be explored again whenever l_3 is explored. This can be avoided by saving l_0 when it is first visited. But the difference from saving l_1 is that saving l_0 is for efficiency and l_1 for termination.

Now we again recall the abstract reachability algorithm in Figure 1 for timed systems. To ensure termination and also to avoid repeated exploration of states (that have more than one predecessors), it saves every new encountered state (l, D) in the passed list when the inclusion-checking for $D \subseteq D'$ fails (i.e. $D \not\subseteq D'$). Obviously this is not necessary if all the predecessors of (l, D) already exist in the PASSED-list. Similar to the case for finite graphs, for termination, we need to save only one state for every *dynamic loop* of a timed automaton.

Definition 4. [Dynamic Loops] Assume a timed automaton with an initial state (l_0, D_0). The set of symbolic states $L_d = \{(l_1, D_1) \ \ldots \ (l_n, D_n)\}$ is a dynamic loop of the timed automaton if $(l_1, D_1) \rightsquigarrow (l_2, D_2) \ \ldots \ (l_{n-1}, D_{n-1}) \rightsquigarrow (l_n, D_n)$ and $(l_n, D_n) \rightsquigarrow (l_1, D_1')$ with $D_1' \subseteq D_1$, and (l_0, D_0) is reachable in the sense that $(l_0, D_0) \rightsquigarrow \ldots \rightsquigarrow (l_1, D_1)$. A symbolic state is said to cover a dynamic loop if it is a member of the loop. □

We claim that to ensure termination, it is sufficient (but not necessary) to save a set of symbolic states that cover all the dynamic loops. Now, the problem is how to compute efficiently such a set.

4.2. Control Structure Analysis and Application

We shall utilize the statical structure of an automaton to identify potential candidates of states to cover dynamic loops.

Definition 5. [Statical Loops and Entry Nodes] A set of nodes $L = \{l_1, \ldots, l_n\}$ of a timed automaton is a statical loop if there is a sequence of edges $l_1 \longrightarrow l_2 \cdots l_{n-1} \longrightarrow l_n$ and $l_n \longrightarrow l_1$ where $l_i \longrightarrow l_j$ denotes that $l_i \xrightarrow{g,r} l_j$ for some g, r is an edge of the automaton. A node $l_i \in L$ is an entry node of the statical loop L if it is an initial node of the automaton or there exists a node $l \notin L$ (outside of the loop) and an edge $l \longrightarrow l_i$. Further, we say that a vector of nodes (i.e. a node of a network) is an entry node if any of its components are entry nodes. □

For example, nodes l_0, l_1, l_2 and l_3 in Figure 5 constitute a statical loop with entry nodes l_0 and l_1; another statical loop is nodes l_1 and l_4 with entry node l_1. In general, since the sets of control nodes and edges of a timed automaton are finite, the number of statical loops is finite and so is the set of entry nodes of all statical loops. In fact the set of entry nodes of a timed automaton can be easily computed by statical analysis using a stack or a slightly modified loop detecting algorithm (see e.g. [29]).

Now note that according to Definition 4, a dynamic loop (a set of symbolic states) must contain a subset of symbolic states whose control nodes constitute a statical loop. As a statical loop always contains an entry node, we have the following fact.

PROPOSITION 1 *Every dynamic loop of a timed automaton contains at least one symbolic state (l, D) where l is an entry node.*

Proof: Standard proof by contradiction. □

Following Proposition 1, to cover all the dynamic loops, we may simply save all the states whose control-nodes are an entry node, and ignore the others. Obviously, this will not give much reduction when dynamic loops include mostly entry nodes, which is the case when a network of automata contains a component whose nodes are mostly entry nodes e.g. a testing automaton. For networks of automata, we adopt the strategy of saving the *first derived* states whose control nodes are an entry node, known as *covering states* in the following sense.

Definition 6. [Covering States] Assume a network of timed automata with an initial state (l_0, D_0) and a given symbolic state (l, D). We say that (l, D) is a covering state of the network if it is reachable in the sense that there exists a sequence of symbolic transitions $(l_0, D_0) \rightsquigarrow (l_1, D_1) \ldots (l_n, D_n) \rightsquigarrow (l, D)$ and

an i (standing for the ith component of the network) such that $l[i]$ is an entry node and $(l_n, D_n) \rightsquigarrow (l, D)$ is derived by an edge $l_n[i] \xrightarrow{g,r} l[i]$ for some g and r. □

From the above definition, it should be obvious that we can easily decide whether a reachable symbolic state is a covering state by an on-the-fly algorithm when the entry nodes of all the component automata are known through statical analysis as discussed earlier.

Finally, we claim that the set of *covering* states of a network covers all its dynamic loops and therefore it suffices to keep them in the passed list for the sake of termination in reachability analysis[7].

THEOREM 3 *Every dynamic loop of a network of timed automata contains at least one covering state.*

Proof: Assume a dynamic loop $L_d = (l_1, D_1) \rightsquigarrow \ldots \rightsquigarrow (l_k, D_k)$ with no covering states. However according to Proposition 1, L_d contains at least one entry node. Further, assume (without loss of generality) that the symbolic state $(l, D) \in L_d$ is an entry node and the components $l[1], \ldots, l[m]$ of l are all in an entry node, and all the other components of l, i.e. $l[m+1], \ldots, l[n]$, are not.

Now, we claim that if L_d contains no covering states, the set of components $l_i[1], \ldots, l_i[m]$ will remain in an entry node in all symbolic states $(l_i, D_i) \in L_d$. Otherwise, if the set of local entry nodes changes, either grows or reduces, it will introduce a covering state. The case of growing is obvious due to the definition for covering states. The argument for the case of reducing is the same as the control nodes of all the components will reach l_1 again by the end of L_d, meaning that the set will sooner or later grows again.

In fact, the assumption that L_d contains no covering states, implies an even stronger property, that is, all symbolic transitions in L_d are derived by components in $l_i[m+1], \ldots, l_i[n]$. A transition is derived by a local transition of a component in $l[1], \ldots, l[m]$, means that the set of local entry nodes will either grow or reduce (discussed above) or the local transition leaves the current entry node and enters an another entry node. The later case implies that the new entry node is a covering state.

Now we construct L'_d by removing $l_i[1], \ldots, l_i[m]$ from all symbolic states $(l_i, D_i) \in L_d$, that is, L'_d contains only the components that are not in an entry nodes. Obviously, all the symbolic transitions of L_d are also in L'_d; thus L'_d must be a loop by definition. However, L'_d contains no components that are in an entry node. This contradicts Proposition 1. □

An improved reachability algorithm according to the saving strategy induced from Theorem 3 (i.e. saving only the covering sates in the passed list) has been implemented in UPPAAL. Our experimental results show that the space-reduction is between 13–72% (see Table 1 and 2 in Section 5).

5. Experimental Results

The techniques developed in preceding sections have been implemented and added to the tool UPPAAL[8][9]. In this section we present the results of an experiment where both the original version of UPPAAL and its extension were applied to verify the following six well-studied examples from the literature:

Philips Audio Protocol (Audio) The protocol was developed and implemented by Philips to exchange control information between components in audio equipment using Manchester encoding. The correctness of the encoding relies on timing delays between signals. It is first studied and manually verified in [10].

We have verified that the main correctness property holds of the protocol, i.e. all bit streams sent by the sender are correctly decoded by the receiver [24], if the timing error is $\pm 5\%$.

Philips Audio Protocol with Bus Collision (Audio w. Collision) This is an extended variant of Philips audio control protocol with bus collision detection [8]. It is significantly larger than the version above since several new components (and variables) are introduced, and existing components are modified to deal with bus collisions.

In the experiment we checked that correct bit sequences are received by the receiver (i.e. Property 1 of [8]), using the error tolerances set by Philips.

Bang & Olufsen Audio/Video Protocol (Bang & Olufsen) This is an audio control protocol highly dependent on real-time. The protocol is developed by Bang & Olufsen, to transmit messages between audio/video components over a single bus, and further studied in [18].

In the experiment we have verified the correctness criteria of the protocol. We refer the reader to Section 5.1 of [18] for more details.

Box Sorter (Box Sorter) The example of [25] is a model of a sorter unit that sorts red and blue boxes. When the boxes moves down a lane they pass a censor and a piston. The sorter reads the information from the censor and sorts out the red boxes by controlling the position of the piston. We have shown, using UPPAAL, that only blue boxes arrive at the end of the lane.

Manufacturing Plant (Manufact. Plant) The example is a model of the manufacturing plant of [28, 14]. It is a production cell with: a 50 feet belt moving from left to right, two boxes, two robots and a service station. Robot A moves boxes off the rightmost extreme of the belt to the service station. Robot B moves boxes from the service station to the left-most extreme of the belt.

Assuming an initial distance between the boxes on the belt we verified that no box will fall off the belt.

Mutual Exclusion Protocol (Mutex 2–Mutex 5) It is the so-called Fischer's protocol that has been studied previously in many experiments, e.g. [1, 30].

	Current	Local		Global		Local+Global	
	#	#	%	#	%	#	%
Audio	828	219	26	774	93	206	25
Audio w. Collision	646 092	198 178	31	370 800	57	111 632	17
Bang & Olufsen	778 288	249 175	32	642 752	83	204 795	26
Box Sorter	625	139	22	175	28	36	6
Manufact. Plant	92 592	27 042	29	50 904	55	14 933	16
Mutex 2	225	44	20	99	44	18	8
Mutex 3	3 376	621	18	1 360	40	240	7
Mutex 4	56 825	9 352	16	22 125	39	3 532	6
Mutex 5	1 082 916	158 875	15	416 556	38	59 720	6
Train Crossing	464	130	28	384	83	114	25

Table 1. Space performance statistics: number of constraints (#) and percentage of Current (%).

The protocol is to ensure mutual exclusion among several processes competing for a critical section using timing constraints and a shared variable. In the experiment we use the version of the protocol where a process may recover from failed attempts to enter the critical section, and also eventually leave the critical section [22].

The protocol is shown to enjoy the invariant property: There is never more than one process existing in the critical section. The results for 2 to 5 processes are shown in Table 1 and 2.

Train Crossing Controller (Train Crossing) It is a variant of the train gate controller [19]. An approaching train signals to the controller which reacts by closing the gate. When the train have passed the controller opens the crossing. We have verified that the gate is closed whenever a train is close to the crossing.

In Table 1 and 2 we present the space (in number of timing constraints stored on the PASSED-buffer) and in the time requirements (in seconds) of the examples on a Sun SPARCstation4 equipped with 64 MB of primary memory. Each example was verified using the current algorithm of UPPAAL (Current), and using modified algorithms for: Compact Data Structure for Constraints (Local), Control Structure Reduction (Global), and their combination (Local+Global).

As shown in Table 1 and 2 both techniques give truly significant space savings: Compact Data Structure for Constraints saves 68–85% of the original consumed

	Current	Local		Global		Local+Global	
	sec	sec	%	sec	%	sec	%
Audio	0.44	0.43	98	0.44	100	0.47	107
Audio w. Collision	3 465.22	2 067.37	60	1 515.88	44	929.22	27
Bang & Olufsen	13 240.49	6 967.38	53	9 348.48	71	4 966.79	38
Box Sorter	0.20	0.18	90	0.41	205	0.41	205
Manufact. Plant	155.61	39.85	26	56.61	36	24.22	16
Mutex 2	0.13	0.14	108	0.15	115	0.14	108
Mutex 3	1.40	0.67	48	0.65	46	0.51	36
Mutex 4	102.49	24.48	24	25.97	25	12.14	12
Mutex 5	14 790.56	3 299.96	22	3 111.21	21	1 138.32	8
Train Crossing	0.19	0.18	95	0.20	105	0.18	95

Table 2. Time performance statistics: seconds (sec) and percentage of Current (%).

space while Control Structure Reduction demonstrates more variation saving 13–72%. Both methods result in better time-performance on the examples consuming more than half a second, whereas the time-perfomance is worse on the smaller examples. Most significant is that the two techniques are completely orthogonal, witnessed by the numbers for the combined technique which shows a space-saving between 75% and 94%.

6. Conclusion

In this paper, we have two contributions to the development of efficient data structures and algorithms for memory-usage reduction in the automated analysis of timed systems.

Firstly, we have presented a compact data structure, for representing the subsets of Euclidean space that arise during verification of timed automata, which provides *minimal* and *canonical* representations for clock constraints, and also allows for efficient inclusion checks between constraint systems. The data structure is based on an $\mathcal{O}(n^3)$ algorithm which, given a constraint systems over real-valued variables consisting of bounds on differences, constructs an equivalent system with a minimal number of constraints. It is essentially a minimization algorithm for weighted directed graphs, that extends the transitive reduction algorithm of [3] to weighted graphs. Given a weighted, directed graph with n vertices, it constructs in time

$\mathcal{O}(n^3)$ a reduced graph with the minimal number of edges having the same shortest path closure as the original graph.

Secondly, we have developed an on-the-fly reduction technique to minimize the space-usage by reducing the total number of symbolic states to save in reachability analysis for timed systems. The technique is based on the observation that to ensure termination in reachability analysis, it is not necessary to save all the explored states in memory, but only certain critical states. Based on static analysis of the control structure of timed automata, we are able to compute a set of *covering states* that cover all the dynamic loops of a system. The set of covering states may not be minimal but sufficient to guarantee termination in an on-the-fly reachability algorithm.

The two techniques and their combination have been implemented in the tool UPPAAL. Our experimental results demonstrate that the techniques result in truly significant space-reductions: For a number of well-studied examples in the literature the space saving is between 75% and 94%, and in all large examples time-performance is improved. Noteworthy is also the observation that the two techniques are completely orthogonal.

As future work, we wish to further study the global on-the-fly reduction technique to identify the *minimal* sets of covering states that ensure termination and also avoid repeated explorations in reachability analysis for timed systems.

Notes

1. Several verification tools for timed systems (e.g. UPPAAL [9]) have been implemented based on this algorithm.
2. For reasons of simplicity and clarity in presentation we have chosen only to consider the non-strict orderings. However, the techniques given extends easily to strict orderings.
3. We assume that D has been simplified to contain at most one upper and lower bound for each clock and clock-difference.
4. To be precise, it is the inclusion between the *solution sets* for D and D'.
5. This would correspond to constraint systems with empty solution set.
6. "$<$" refers to the assumed ordering on the vertices of G.
7. Note that this is only a sufficient condition but not necessary.
8. For more information about the tool UPPAAL, see the web site `http://www.uppaal.com/`.

References

1. Martin Abadi and Leslie Lamport. An Old-Fashioned Recipe for Real Time. In *Proc. of REX Workshop "Real-Time: Theory in Practice"*, number 600 in Lecture Notes in Computer Science, 1992.
2. Luca Aceto, Auguto Bergueno, and Kim G. Larsen. Model checking via reachability testing for timed automata. In Bernard Steffen, editor, *Proc. of the 4th Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, number 1384 in Lecture Notes in Computer Science, pages 263–280. Springer–Verlag, 1998.
3. A.V. Aho, M.R. Garey, and J.D. Ullman. The Transitive Reduction of a Directed Graph. *SIAM Journal on Computing*, 1(2):131–137, June 1972.

4. R. Alur and D. Dill. Automata for Modelling Real-Time Systems. In *Proc. of Int. Colloquium on Algorithms, Languages and Programming*, number 443 in Lecture Notes in Computer Science, pages 322–335, July 1990.
5. H. R. Andersen. Partial Model Checking. In *Proc. of Symp. on Logic in Computer Science*, 1995.
6. Eugene Asarin, Oded Maler, and Amir Pnueli. Data–structures for the verification of timed automata. In *Proc. of the Int. Workshop on Hybrid and Real-Time Systems*, 1997.
7. Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957.
8. Johan Bengtsson, W.O. David Griffioen, Kare J. Kristoffersen, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, and Wang Yi. Verification of an Audio Protocol with Bus Collision Using UPPAAL. In Rajeev Alur and Thomas A. Henzinger, editors, *Proc. of the 8*th *Int. Conf. on Computer Aided Verification*, number 1102 in Lecture Notes in Computer Science, pages 244–256. Springer–Verlag, July 1996.
9. Johan Bengtsson, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, and Wang Yi. UPPAAL in 1995. In *Proc. of the 2*nd *Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, number 1055 in Lecture Notes in Computer Science, pages 431–434. Springer–Verlag, March 1996.
10. D. Bosscher, I. Polak, and F. Vaandrager. Verification of an Audio-Control Protocol. In *Proc. of Formal Techniques in Real-Time and Fault-Tolerant Systems*, number 863 in Lecture Notes in Computer Science, 1994.
11. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic Model Checking: 10^{20} states and beyond. In *Proc. of IEEE Symp. on Logic in Computer Science*, 1990.
12. E. M. Clarke, T. Filkorn, and S. Jha. Exploiting Symmetry in Temporal Logic Model Checking. In *Proc. of the 5*th *Int. Conf. on Computer Aided Verification*, number 697 in Lecture Notes in Computer Science, 1993.
13. E. M. Clarke, O. Grümberg, and D. E. Long. Model Checking and Abstraction. *Principles of Programming Languages*, 1992.
14. C. Daws and S. Yovine. Two examples of verification of multirate timed automata with KRONOS. In *Proc. of the 16*th *IEEE Real-Time Systems Symposium*, pages 66–75, December 1995.
15. David Dill. Timing assumptions and verification of finite-state concurrent systems. In J. Sifakis, editor, *Proc. of Automatic Verification Methods for Finite State Systems*, number 407 in Lecture Notes in Computer Science, pages 197–212. Springer–Verlag, 1989.
16. E. A. Emerson and C. S. Jutla. Symmetry and Model Checking. In *Proc. of the 5*th *Int. Conf. on Computer Aided Verification*, number 697 in Lecture Notes in Computer Science, 1993.
17. P. Godefroid and P. Wolper. A Partial Approach to Model Checking. In *Proc. of IEEE Symp. on Logic in Computer Science*, pages 406–415, 1991.
18. Klaus Havelund, Arne Skou, Kim G. Larsen, and Kristian Lund. Formal Modeling and Analysis of an Audio/Video Protocol: An Industrial Case Study Using UPPAAL. In *Proc. of the 18*th *IEEE Real-Time Systems Symposium*, December 1997.
19. Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. A Users Guide to HYTECH. Technical report, Department of Computer Science, Cornell University, 1995.
20. Thomas. A. Henzinger, Xavier Nicollin, Joseph Sifakis, and Sergio Yovine. Symbolic Model Checking for Real-Time Systems. *Information and Computation*, 111(2):193–244, 1994.
21. Gerard Holzmann. *The Design and Validation of Computer Protocols*. Prentice Hall, 1991.
22. Kare J. Kristoffersen, Francois Larroussinie, Kim G. Larsen, Paul Pettersson, and Wang Yi. A Compositional Proof of a Real-Time Mutual Exclusion Protocol. In *Proc. of the 7*th *Int. Joint Conf. on the Theory and Practice of Software Development*, April 1997.
23. Kim G. Larsen, Paul Pettersson, and Wang Yi. Compositional and Symbolic Model-Checking of Real-Time Systems. In *Proc. of the 16*th *IEEE Real-Time Systems Symposium*, pages 76–87, December 1995.
24. Kim G. Larsen, Paul Pettersson, and Wang Yi. Diagnostic Model-Checking for Real-Time Systems. In *Proc. of Workshop on Verification and Control of Hybrid Systems III*, number 1066 in Lecture Notes in Computer Science, pages 575–586. Springer–Verlag, October 1995.

25. Kim G. Larsen, Paul Pettersson, and Wang Yi. UPPAAL in a Nutshell. *Int. Journal on Software Tools for Technology Transfer*, 1(1–2):134–152, October 1997.
26. Florence Pagani. Partial orders and verification of real-time systems. In Bengt Jonsson and Joachim Parrow, editors, *Proc. of Formal Techniques in Real-Time and Fault-Tolerant Systems*, number 1135 in Lecture Notes in Computer Science, pages 327–346. Springer–Verlag, 1996.
27. Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
28. A. Puri and P. Varaiya. Verification of hybrid systems using abstractions. In *Hybrid Systems Workshop*, number 818 in Lecture Notes in Computer Science. Springer–Verlag, October 1994.
29. Robert Sedgewick. *Algorithms*. Addison-Wesley, 2*nd* edition, 1988.
30. N. Shankar. Verification of Real-Time Systems Using PVS. In *Proc. of the* 5th *Int. Conf. on Computer Aided Verification*, number 697 in Lecture Notes in Computer Science. Springer–Verlag, 1993.
31. A. Valmari. A Stubborn Attack on State Explosion. *Theoretical Computer Science*, 3, 1990.
32. M.Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. of Symp. on Logic in Computer Science*, pages 322–331, June 1986.
33. Mihalis Yannakakis and David Lee. An efficient algorithm for minimizing real–time transition systems. In *Proc. of the* 5th *Int. Conf. on Computer Aided Verification*, number 697 in Lecture Notes in Computer Science, pages 210–224, 1993.
34. Wang Yi, Paul Pettersson, and Mats Daniels. Automatic Verification of Real-Time Communicating Systems By Constraint-Solving. In *Proc. of the* 7th *Int. Conf. on Formal Description Techniques*, pages 223–238, 1994.