

Verics 2006 - a Model Checker for Real-Time and Multi-Agent Systems*

Magdalena Kacprzak¹, Wojciech Nabiałek², Artur Niewiadomski²,
Wojciech Penczek^{2,3}, Agata Pólrola⁴, Maciej Szreter³,
Bożena Woźna⁵, and Andrzej Zbrzezny⁵

¹ Białystok University of Technology, FCS, Wiejska 45a, 15-351 Białystok, Poland

² Podlasie Academy, ICS, Sienkiewicza 51, 08-110 Siedlce, Poland

³ ICS PAS, Ordonia 21, 01-237 Warsaw, Poland

⁴ University of Lodz, FMCS, Banacha 22, 90-238 Lodz, Poland

⁵ Jan Długosz Academy, IMCS, Armii Krajowej 13/15, 42-200 Częstochowa, Poland
verics@ipipan.waw.pl

Abstract. The papers presents the current stage of the development of VerICS - a model checker for real-time and multi-agent systems. Depending on the type of a system considered, it enables to test various classes of properties - from reachability to temporal, epistemic and deontic formulas. The model checking methods used to this aim include both SAT-based and enumerative ones. In the paper we focus on new features of the verifier: SAT-based model checking for multi-agent systems and several extensions and improvements to real-time systems' verification.

1 Introduction

The paper presents the current stage of the development of VerICS, a model checker for real-time and multi-agent systems. Depending on the type of a system considered, the verifier enables to test various classes of properties - from reachability of a state satisfying certain conditions to more complicated features expressed by formulas of (timed) temporal, epistemic, or deontic logics. The model checking methods implemented include both SAT-based and enumerative ones (where by the latter we mean these consisting in generating abstract models for systems). Our first work [3] presenting VerICS dealt mainly with verification of Real-Time Systems (RTS). In this paper we focus on VerICS' new features, i.e., SAT-based model checking for MAS (Multi-Agent Systems), and several extensions to RTS verification.

2 Related Work

VerICS enables verification of real-time and multi-agent systems. A survey of model checkers for RTS can be found in [16]. Considering MAS verification,

* Partly supported by the Ministry of Education and Science under the grant No. 3 T11C 011 28

Verics is, to our best knowledge, one of the three existing model checkers for verifying MAS directly, and the only one which applies SAT to this aim. The other two: MCMAS [11] and MCK [6] implement BDD-based verification methods. Some other tools like CASP [1, 20] or MABLE [20] enable translations from MAS to languages accepted by “general purpose” model checkers like Spin or JavaPathFinder.

3 Theoretical Background

A network of communicating (timed) automata is the basic Verics’ formalism for modelling a system to be verified. Timed automata are used to specify RTS (possibly with clock differences expressing constraints on their behaviour), whereas timed or untimed automata are applied to model MAS (possibly extended in a way to handle certain features of interest, like deontic automata in [9]).

The tuples of *local states* of the automata in a network N define the *global states* of the system considered. The set of all the possible *runs* (i.e., infinite evolutions from a given initial state) of an RTS modelled by N gives us a *computation tree* which, after labelling the states with propositions from a given set PV which are true at these states (i.e., changing the tree into a *model*), is used to interpret the formulas of timed or untimed temporal logics (like CTL or TCTL) expressing properties to be checked. In the case of modelling a MAS we augment the model with *epistemic* or *deontic accessibility relations*. The resulting structure enables us to interpret formulas involving temporal operators, epistemic operators - to reason about what agents *know* [5], and deontic operators - to reason about *correctness* of their behaviour.

SAT-based verification methods represent the models and properties of systems in the form of boolean formulas in order to reduce the state explosion. These for MAS involve bounded (BMC) and unbounded model checking (UMC). Currently, Verics implements UMC for CTL_pK (Computation Tree Logic with knowledge and past operators) [8], and BMC for ECTLKD (the existential fragment of CTL extended with knowledge and deontic operators) [9, 15, 21, 22] as well as TECTLK (the existential fragment of timed CTL extended with knowledge operators) [12]. Considering verification of RTS, it offers BMC for proving (un)reachability [23] (also for timed automata with clock differences [24]), and UMC for proving CTL properties for slightly restricted timed automata [18]. Below we present some more details of BMC for ECTLK and UMC for CTL_pK .

3.1 Bounded Model Checking

Bounded Model Checking (BMC) is a symbolic method aimed at verification of temporal properties of distributed (timed) systems. It is based on the observation that some properties of a system can be checked over a part of its model only. In the simplest case of reachability analysis, this approach consists in an iterative encoding of a finite symbolic path (computation) as a propositional formula.

The satisfiability of the resulting propositional formula is then checked using an external SAT-solver.

Consider a system consisting of n agents. Let G be the set of global states (i.e., tuples of n local states) of the system under consideration, g^0 be its initial state, $Q \subseteq G$ be the set of states reachable from g^0 , $T \subseteq G \times G$ be a transition relation, $\sim_i \subseteq G \times G$ for $i = 1, \dots, n$, be the epistemic accessibility relation¹, and $V : G \rightarrow 2^{PV}$ be a valuation function. Moreover, let for $k \in \mathbb{N}_+$ a k -path be a finite sequence of $k + 1$ states $\pi = (g_0, \dots, g_k)$, where $g_i \in G$ for $i = 0, \dots, k$ and $(g_i, g_{i+1}) \in T$ for each $0 \leq i < k$. For a k -path $\pi = (g_0, \dots, g_k)$, let $\pi(i) = g_i$ for each $0 \leq i \leq k$. By $\Pi_k(g)$ we denote the set of all the k -paths starting at g . We assume that the reader is familiar with the standard syntax and semantics of CTL, CTLK, and CTL_PK (details can be found in [10]). In order to restrict the semantics of ECTLK to a part of the model we introduce the following definition:

Definition 1. Let $M = (G, Q, g^0, T, \sim_1, \dots, \sim_n, V)$ be a model and $k \in \mathbb{N}_+$. The k -model for M is defined as $M_k = (Q, g^0, P_k, \sim_1, \dots, \sim_n, V')$, where P_k is the set of all the k -paths of M over Q , i.e., $P_k = \bigcup_{s \in Q} \Pi_k(s)$, and $V' = V|_Q$.

In order to identify k -paths that represent infinite paths we define the function $loop : P_k \rightarrow 2^{\mathbb{N}}$ as: $loop(\pi) = \{l \mid 0 \leq l \leq k \text{ and } (\pi(k), \pi(l)) \in T\}$, which returns the set of indices l of π for which there is a transition from $\pi(k)$ to $\pi(l)$. Then, we define a *bounded semantics* of ECTLK (we omit here the operators \overline{D} , \overline{E} and \overline{C}):

Definition 2. Let M_k be a k -model and α, β be ECTLK formulas. $M_k, s \models \alpha$ denotes that α is true at the state s of M_k . M_k is omitted if it is clear from the context. The relation \models for modal operators is defined inductively as follows:

$$\begin{aligned} s \models EX\alpha & \quad \text{iff } (\exists \pi \in P_k(s)) \pi(1) \models \alpha, \\ s \models EG\alpha & \quad \text{iff } (\exists \pi \in P_k(s)) (\forall 0 \leq j \leq k) (\pi(j) \models \alpha \text{ and } loop(\pi) \neq \emptyset), \\ s \models E(\alpha U \beta) & \quad \text{iff } (\exists \pi \in P_k(s)) (\exists 0 \leq j \leq k) (\pi(j) \models \beta \text{ and } (\forall 0 \leq i < j) \pi(i) \models \alpha), \\ s \models \overline{K}_i \alpha & \quad \text{iff } (\exists \pi \in P_k(s)) (\exists 0 \leq j \leq k) (\pi(j) \models \alpha \text{ and } s \sim_i \pi(j)), \end{aligned}$$

Model checking over models can be reduced to model checking over k -models. The main idea of BMC for ECTLK is that we can check φ over M_k by testing the satisfiability of the propositional formula $[M, \varphi]_k := [M^{\varphi, g^0}]_k \wedge [\varphi]_{M_k}$, where the first conjunct represents (a part of) the model under consideration and the second a number of constraints that must be satisfied on M_k for φ to be satisfied. Once this translation is defined, checking satisfiability of an ECTLK formula can be done by means of a SAT-checker. Typically, we start with $k := 1$, test satisfiability for the translation, and increase k by one until either $[M^{\varphi, g^0}]_k \wedge [\varphi]_{M_k}$ becomes satisfiable, or k reaches the maximal depth of M .²

Translation. Below, we provide some details of the translation. Given a MAS represented by a network of automata. Each global state g of the system can be

¹ Two global states are in the relation \sim_i if they share the same i -th local states.

² The upper limit is $|Q|$.

represented by $w = (w[1], \dots, w[m])$, for some $m \in \mathbb{N}_+$, (which we shall call a *global state variable*), where each $w[i]$ for $i = 1, \dots, m$ is a propositional variable. A sequence $w_{0,j}, \dots, w_{k,j}$ of global state variables is called a *symbolic k -path j* .

Let f_k be a function which determines the number of k -paths sufficient for checking a given ECTLK formula (see [21]). Moreover, let $\overline{PV} = \{\neg p \mid p \in PV\}$ be the set of negated propositions of PV , and let $lit: \{0, 1\} \times PV \rightarrow PV \cup \overline{PV}$ be a function defined as follows: $lit(0, p) = \neg p$ and $lit(1, p) = p$. Furthermore, let w, v be two global state variables. We use the formula $I_g(w) := \bigwedge_{i=1}^m lit(g_i, w[i])$ to encode a global state $g = (g_1, \dots, g_m)$ of the model, i.e., if $g_i = 1$, then it is encoded by $w[i]$, and if $g_i = 0$, then it is encoded by $\neg w[i]$. The propositional formula $[M^{\varphi, g^0}]_k$, representing the k -paths in the k -model, is defined as $[M^{\varphi, g^0}]_k := I_{g^0}(w_{0,0}) \wedge \bigwedge_{j=1}^{f_k(\varphi)} \bigwedge_{i=0}^{k-1} T(w_{i,j}, w_{i+1,j})$, where $w_{0,0}$ and $w_{i,j}$ for $0 \leq i \leq k$ and $1 \leq j \leq f_k(\varphi)$ are global state variables, and $T(w_{i,j}, w_{i+1,j})$ is a formula encoding the transition relation T .

The next step of the algorithm consists in translating an ECTLK formula φ into a propositional formula. Let w, v be global state variables. We make use of the following propositional formulas in the encoding:

- $p(w)$ encodes a proposition p of ECTLK over w ;
- $H(w, v) := \bigwedge_{j=1}^m w[j] \Leftrightarrow v[j]$ represents logical equivalence between global state encodings u and v (i.e., encodes that u and v represent the same global states);
- $HK_i(w, v) := \bigwedge_{j \in I_{x_i}} w[j] \Leftrightarrow v[j]$ represents logical equivalence between i -local state encodings u and v , (i.e., encodes that u and v share i -local states);
- $L_{k,j}(l)$ encodes a backward loop connecting the k -th state to the l -th state in the symbolic k -path j , for $0 \leq l \leq k$.

The translation of φ at the state $w_{m,n}$ into the propositional formula $[\varphi]_k^{[m,n]}$ is as follows (we give the translation of selected formulas only):

$$\begin{aligned} [\text{EX}\alpha]_k^{[m,n]} &:= \bigvee_{i=1}^{f_k(\varphi)} \left(H(w_{m,n}, w_{0,i}) \wedge [\alpha]_k^{[1,i]} \right), \\ [\text{EG}\alpha]_k^{[m,n]} &:= \bigvee_{i=1}^{f_k(\varphi)} \left(H(w_{m,n}, w_{0,i}) \wedge \left(\bigvee_{l=0}^k L_{k,i}(l) \right) \wedge \bigwedge_{j=0}^k [\alpha]_k^{[j,i]} \right), \\ [\text{E}(\alpha \cup \beta)]_k^{[m,n]} &:= \bigvee_{i=1}^{f_k(\varphi)} \left(H(w_{m,n}, w_{0,i}) \wedge \bigvee_{j=0}^k \left([\beta]_k^{[j,i]} \wedge \bigwedge_{t=0}^{j-1} [\alpha]_k^{[t,i]} \right) \right), \\ [\overline{\text{K}}_l \alpha]_k^{[m,n]} &:= \bigvee_{i=1}^{f_k(\varphi)} \left(I_{g^0}(w_{0,i}) \wedge \bigvee_{j=0}^k \left([\alpha]_k^{[j,i]} \wedge HK_l(w_{m,n}, w_{j,i}) \right) \right). \end{aligned}$$

Given the translations above, we can now check φ over M_k by checking the satisfiability of the propositional formula $[M^{\varphi, g^0}]_k \wedge [\varphi]_{M_k}$, where $[\varphi]_{M_k} = [\varphi]_k^{[0,0]}$.

3.2 Unbounded Model Checking

Unlike BMC, UMC is capable of handling the whole language of the logic. Our aim is to translate CTL_pK formulas into propositional formulas in conjunctive normal form. Specifically, for a given CTL_pK formula φ we compute a corresponding propositional formula $[\varphi](w)$, where w is a global state variable (i.e., a vector of propositional variables for representing global states) encoding these states of the model where φ holds. To calculate the actual translations we use

either the QBF or the fixed-point characterisation of CTL_pK formulas. Below, we recall the necessary notions.

Formulas in CNF and QBF. Let PV be a finite set of propositional variables. A *literal* is a propositional variable $p \in PV$ or its negation $\neg p$. A *clause* is a disjunction of a set of zero or more literals $l[1] \vee \dots \vee l[n]$. A disjunction of zero literals is taken to mean the constant **false**. A formula is in a *conjunctive normal form* (CNF) if it is a conjunction of a set of zero or more clauses $c[1] \wedge \dots \wedge c[n]$. An *assignment* A is a partial function from PV to $\{\mathbf{true}, \mathbf{false}\}$.

In our method, in order to have a more succinct notation for complex operations on boolean formulas, we also use *Quantified Boolean Formulas* (QBF), an extension of propositional logic by means of quantifiers ranging over propositions. The BNF syntax of a QBF formula is given by:

$$\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \exists p.\alpha \mid \forall p.\alpha.$$

The semantics of the quantifiers is defined by $\exists p.\alpha$ iff $\alpha(p \leftarrow \mathbf{true}) \vee \alpha(p \leftarrow \mathbf{false})$, and $\forall p.\alpha$ iff $\alpha(p \leftarrow \mathbf{true}) \wedge \alpha(p \leftarrow \mathbf{false})$, where $\alpha \in \text{QBF}$, $p \in PV$ and $\alpha(p \leftarrow q)$ denotes substitution with the variable q of every occurrence of the variable p in formula α . We will use the notation $\forall v.\alpha$, where $v = (v[1], \dots, v[m])$ is a vector of propositional variables, to denote $\forall v[1].\forall v[2].\dots.\forall v[m].\alpha$. For a given QBF formula $\forall v.\alpha$, we can construct a CNF formula equivalent to it by using the algorithm *forall* [13].

Our aim is to translate the whole CTL_pK language into boolean formulas. To this aim, we first need to translate CTL_pK formulas into QBF form. Before doing so, we need to be able to present the CTL_pK language in terms of fixed-points. This is shown in the next subsection. In our presentation, we follow and adapt definitions given in [2].

Fixed-point characterisation of CTL_pK . Let $M = (G, Q, g^0, T, \sim_1, \dots, \sim_n, V)$ be a model defined as in Sec. 3.1. Notice that the set 2^G of all subsets of G forms a lattice under the set inclusion ordering. Each element G' of the lattice can also be thought of as a *predicate* on G , where the predicate is viewed as being true for exactly the states in G' . The least element in the lattice is the empty set, which corresponds to the predicate **false**, and the greatest element in the lattice is the set G , which corresponds to **true**. A function τ mapping 2^G to 2^G is called a *predicate transformer*. A set $G' \subseteq G$ is a *fixed point* of a function $\tau : 2^G \rightarrow 2^G$ if $\tau(G') = G'$.

Whenever τ is monotonic (i.e., when $P \subseteq P'$ implies $\tau(P) \subseteq \tau(P')$), a function τ has a least fixed point denoted by $\mu Z.\tau(Z)$, and a greatest fixed point, denoted by $\nu Z.\tau(Z)$. When τ is monotonic and \bigcup -continuous (i.e., when $P_1 \subseteq P_2 \subseteq \dots$ implies $\tau(\bigcup_i P_i) = \bigcup_i \tau(P_i)$), then $\mu Z.\tau(Z) = \bigcup_{i \geq 0} \tau^i(\mathbf{false})$. When τ is monotonic and \bigcap -continuous (i.e., when $P_1 \supseteq P_2 \supseteq \dots$ implies $\tau(\bigcap_i P_i) = \bigcap_i \tau(P_i)$), then $\nu Z.\tau(Z) = \bigcap_{i \geq 0} \tau^i(\mathbf{true})$ (see [19]).

In order to obtain fixed-point characterisations of the modal operators, we identify each CTL_pK formula α with the set $\langle \alpha \rangle_M$ of the states in M at which

this formula is true, formally $\langle \alpha \rangle_M = \{s \in G \mid M, s \models \alpha\}$. If M is clear from the context we omit the subscript M. Furthermore, we define functions AX, AY, K_i from 2^G to 2^G as follows:

- AX(Z) = $\{g \in G \mid \forall g' \in G \text{ if } (g, g') \in T, \text{ then } s' \in Z\}$,
- AY(Z) = $\{g \in G \mid \forall g' \in G \text{ if } (g', g) \in T, \text{ then } s' \in Z\}$,
- $K_i(Z)$ = $\{g \in G \mid \forall g' \in G \text{ if } (g^0, g') \in T^* \text{ and } g \sim g', \text{ then } g' \in Z\}$,

Observe that $\langle O\alpha \rangle = O(\langle \alpha \rangle)$, for $O \in \{\text{AX}, \text{AY}, K_i\}$. Then, the following temporal and epistemic operators may be characterised as the least or the greatest fixed point of an appropriate monotonic (\cap -continuous or \cup -continuous) predicate transformer (see [4, 2]): $\langle \text{AG}\alpha \rangle = \nu Z. \langle \alpha \rangle \cap \text{AX}(Z)$, $\langle \text{A}(\alpha \text{U}\beta) \rangle = \mu Z. \langle \beta \rangle \cup (\langle \alpha \rangle \cap \text{AX}(Z))$, $\langle \text{AH}\alpha \rangle = \nu Z. \langle \alpha \rangle \cap \text{AY}(Z)$.

Unbounded Model Checking on CTL_pK . Given a model M , we encode its states like for BMC in Sec. 3.1. Now, our aim is to translate CTL_pK formulas into propositional formulas. Specifically, for a given CTL_pK formula β we compute a corresponding propositional formula $[\beta](w)$, which encodes those states of the system that satisfy the formula. Operationally, we work outwards from the most nested subformulas, i.e., the atoms. In other words, to compute $[O\alpha](w)$, where O is a modality, we work under the assumption of already having computed $[\alpha](w)$. To calculate the actual translations we use either the fixed-point or the QBF characterisation of CTL_pK formulas. For example, the formula $[\text{AX}\alpha](w)$ is equivalent to the QBF formula $\forall v. (T(w, v) \Rightarrow [\alpha](v))$. We can use similar equivalences for formulas $\text{AY}\alpha, K_i\alpha$. More specifically, we use three basic algorithms: The first one, implemented by the procedure *forall* [13], is used for formulas $O\alpha$ such that $O \in \{\text{AX}, \text{AY}, K_i\}$. This procedure eliminates the universal quantifier from a QBF formula representing a CTL_pK formula, and returns the result in a conjunctive normal form. The second algorithm [8], implemented by the procedure *gfpO*, is applied to formulas $O\alpha$ such that $O \in \{\text{AG}, \text{AH}\}$. This procedure computes the greatest fixed point. For formulas of the form $\text{A}(\alpha \text{U}\beta)$ we use the third procedure, called *lfpAU*, which computes the least fixed point. In so doing, given a formula β we obtain a propositional formula $[\beta](w)$ such that β is valid in the model M iff the propositional formula $[\beta](w) \wedge I_{g^0}(w)$ is satisfiable.

4 Implementation

The architecture of Verics is shown in Fig. 1. The system consists of:

- **Estelle to Intermediate Language (IL) translator**, which enables to handle specifications written in a subset of Estelle [7] (the standardised language for specifying communicating protocols and distributed systems);
- **IL to timed automata translator**, which, given an IL specification, generates the corresponding network of timed automata or the global timed automaton;

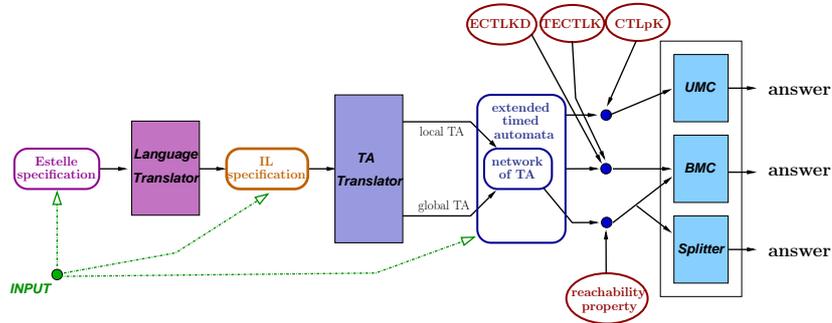


Fig. 1. Architecture of VerICS

- **BMC module**, which implements BMC-based verification for the classes of properties shown in the figure. The SAT-solver used is MiniSat [14] or RSat [17]; the system can be configured to work with other solvers;
- **UMC module**, which provides preliminary implementations of UMC verification methods for properties described above. The module is integrated with a modified version of the SAT-solver ZChaff [25];
- **Splitter module**, which performs reachability verification on abstract models generated for timed automata.

Verics has been implemented in C++; its internal functionalities are available via a interface written in Java. The current distribution (binaries of a standalone program to be run under Linux, or a client version which can work under an arbitrary Java-supporting operating system) can be downloaded from <http://verics.ipipan.waw.pl>.

5 Experimental Results

One of the important elements taken into account while rating a model checker is the efficiency of its behaviour. In this section we present some examples of multi-agent and real-time systems we tested, and provide the experimental results obtained. All the examples are standard (scalable) benchmarks, used typically to test efficiency of low-level verification modules. The results we obtain for them allow to compare the effectiveness of Verics with the other tools available.

The examples we tested were:

- **Dining Cryptographers** - a system consisting of n agents - cryptographers having dinner in a restaurant, who want to know who paid for their meal: one of them or the agency. To reach that while keeping the payer anonymous, each agent flips the coin and states aloud whether the outcomes obtained by him and by the neighbour on his right are equal or different, saying the opposite to what he sees if he has paid. An odd number of differences uttered

indicates that a cryptographer was the payer; an even number means that the one who has paid was the agency.

In order to show the correctness of the above protocol, we prove that the formulas specifying its desired properties hold, while these expressing the undesired ones are false. In our tests we deal with the following formulas:

$$\begin{aligned}\varphi_D^1 &:= \text{AG}(\text{odd} \wedge \neg \text{paid}_1 \Rightarrow \text{K}_1(\bigvee_{i=2,\dots,n} \text{paid}_i)), \\ \varphi_D^2 &:= \text{AG}(\text{odd} \wedge \neg \text{paid}_1 \Rightarrow \bigvee_{i=2,\dots,n} \text{K}_1(\text{paid}_i)), \\ \varphi_D^3 &:= \text{AG}(\neg \text{paid}_1 \Rightarrow \text{K}_1(\bigvee_{i=2,\dots,n} \text{paid}_i)).\end{aligned}$$

The formula φ_D^1 expresses that always when the number of statements “different” is odd and the first cryptographer has not paid for the dinner, then he knows that another cryptographer paid. The formula is proven to hold in the system. The next formula φ_D^2 , which violates the requirements towards the protocol, says that always when the number of the statements “different” is odd and the first cryptographer has not paid for dinner, then he knows the cryptographer who has paid. The formula is obviously false, since the above information should be secret - none of the cryptographers should know the payer, and an odd number of differences means only that the one who has paid was not the agency. The third formula specifies the property which means that always when the first cryptographer has not paid for dinner, then he knows that some other cryptographer has paid. This, again, is a false property, as the agency can be the payer as well.

The results for the above formulas are presented in Fig. 2. Analysing the results one can notice that the verification techniques applied can be seen as complementary: the UMC method gives better results for the formulas φ_D^1 and φ_D^2 , while for φ_D^3 bounded model checking is far more efficient. In the latter case verifying such a big system (up to 1000 cryptographers!) is possible due to the fact that a counterexample can be found on a path which is relatively short. For the first two formulas the length of the path should be incremented up to the model’s diameter. In addition to that, the formulas have to be tested on several symbolic paths, which influences the size of the model we are able to test this way.

formula	n	BMC [s/MB]	SAT [s/MB]	n	UMC [s/MB]
φ_D^1	4	0.8 / 9.2	81224 / *	17	367.0 / *
φ_D^2	4	15.97 / 25.8	9359.24 / *	9	392.0 / *
φ_D^3	1000	329 / 1885.0	29.67 / *	15	421.0 / *

Fig. 2. Experimental results (the star denotes that no data are available)

k	BMC		MiniSat		
	s	MB	s	MB	satisfiable
0	0.2	4.6	0.0	2.9	no
2	1.0	8.5	0.8	9.8	no
5	2.0	13.8	27.7	37.3	no
8	3.2	19.2	1270.8	521.1	no
11	4.3	24.6	2408.3	852.1	no
14	5.6	29.9	4267.4	1527.7	no
17	7.0	35.3	590.8	419.4	YES
Total	23.3	35.3	8565.8	1527.7	

Fig. 3. Experimental results for the Fischer protocol with the parameters which make the mutual exclusion violated ($\Delta = 2$, $\delta = 1$). The number of processes involved is 80

- **Fischer’s Mutual Exclusion Protocol** - a system consisting of n processes trying to enter their critical sections, and a process which coordinates their access. The behaviour of the system depends on the values δ and Δ ($\delta < \Delta$ makes it incorrect - the mutual exclusion does not hold). The property we have tested is expressed by the formula

$$\varphi_M := \text{EF} \left(\bigvee_{i,j=1,\dots,n;i \neq j} (\text{crit}_i \wedge \text{crit}_j) \right),$$

which states that two different processes can be in their critical sections at the same time (i.e., the mutual exclusion is violated).

The tables in Figures 3 and 4 show the results for verifying the above protocol using the bounded model checking method. In Fig. 3, a counterexample proving that the mutual exclusion does not hold, is found on the path of the length 17. The table in Fig. 4 presents the process of testing unreachability of the property in a system which satisfies the requirements. To this aim, it is checked alternately whether in the system there exists a *free* path [23] of a given length, and whether the property holds on a path of such a length which starts at the initial state of the system. Then, if necessary, the length of the path is incremented. In the case considered in the table it occurred that no free path of length 53 can be found, whereas on all the shorter paths the property tested does not hold. This implies that the property does not hold for the system at all.

Fig. 5 compares the results obtained using the UMC method provided by VerICS with these obtained using some other model checkers available. The table shows that in most the cases VerICS occurred to be more efficient.

- **Timed Alternating Bit Protocol** - a system consisting of two agents: a sender and a receiver, which exchange messages over two unreliable communication channels, choosing the channel to be used on the basis of the round trip time of a control bit. The properties we have tested were “for each execution of the protocol, if the receiver got an acknowledgement in the time not exceeding t_1 and the value of the bit sent was 0 then the receiver knows

k	BMC			MiniSat		
	free path	s	MB	s	MB	satisfiable
0	-	0.0	2.8	0.0	1.7	no
2	+	0.0	3.2	0.0	2.0	yes
2	-	0.0	3.2	0.0	2.1	no
...						
20	+	0.2	6.8	1.6	5.6	yes
20	-	0.2	6.8	5.9	8.8	no
...						
50	+	0.6	12.6	2670.9	206.1	yes
50	-	0.5	12.6	148.1	71.6	no
53	+	0.6	13.2	3992.1	233.1	NO
Total		10.4	13.2	10037.9	233.1	

Fig. 4. Experimental results for the Fischer protocol with the parameters for which the mutual exclusion holds ($\Delta = 1$, $\delta = 2$). The number of processes involved is 10

Parameters	Time [s]		
	UppAal	RED	VerICS UMC
$N = 10, \Delta = 1, \delta = 2$	37	53	34
$N = 11, \Delta = 1, \delta = 2$	121	141	46
$N = 12, \Delta = 1, \delta = 2$	580	304	59
$N = 13, \Delta = 1, \delta = 2$	-	657	88
$N = 15, \Delta = 1, \delta = 2$	-	-	154
$N = 18, \Delta = 1, \delta = 2$	-	-	376
$N = 20, \Delta = 1, \delta = 2$	-	-	491
$N = 10, \Delta = 3, \delta = 4$	33	53	50
$N = 11, \Delta = 3, \delta = 4$	125	133	71
$N = 10, \Delta = 2, \delta = 1$	7	49	97

Fig. 5. A comparison of the results for the Fischer protocol, obtained using VerICS' UMC and the tools UppAal and RED

the value of this bit” and “for each execution of the protocol if the receiver got an acknowledgement in the time not exceeding t_1 and the value of the bit sent was 0 then the sender knows that the receiver knows the value of this bit”, which are expressed, respectively, by the following two formulas:

$$\varphi_A^1 := \text{AG}_{[0, t_1]}((\text{recack} \wedge \text{bit}0) \Rightarrow \text{K}_R(\text{bit}0)),$$

$$\varphi_A^2 := \text{AG}_{[0, t_1]}((\text{recack} \wedge \text{bit}0) \Rightarrow \text{K}_S \text{K}_R(\text{bit}0)).$$

The results of searching for counterexamples for these properties, obtained using the BMC module of VerICS aimed at TECTLK verification, are presented in Fig. 6. In contrast to other benchmarks tested we provide no comparison with other tools, as, to our knowledge, there are no other model checkers which enable TECTLK verification.

k	$\neg\varphi_A^1$					$\neg\varphi_A^2$				
	BMC		MiniSat			BMC		MiniSat		
	s	MB	s	MB	satisfiable	s	MB	s	MB	satisfiable
0	0.0	1.4	0.0	3.6	no	0.0	1.4	0.0	3.6	no
1	0.0	1.8	0.0	3.9	no	0.0	1.8	0.0	4.0	no
2	0.0	2.3	0.0	4.4	no	0.1	2.3	0.0	4.4	no
...										
9	0.9	10.1	0.3	11.2	no	5.6	10.3	0.2	11.0	no
10	1.2	12.7	0.3	13.1	no	8.6	12.7	0.4	13.3	no
11	1.6	14.8	1.6	15.4	yes	13.0	14.8	1.4	15.4	yes
Total	5.6	14.8	2.6	15.4		35.0	14.8	2.4	15.4	

Fig. 6. Experimental results for the timed alternating bit protocol

6 Final Remarks

In the paper, we presented an overview of the model checker VerICS, focusing on its new features (comparing with the version presented in [3]), i.e., SAT-based verification of multi-agent systems and some extensions and improvements to real-time systems' verification. VerICS offers several capabilities which, to our knowledge, are available in no other tools. It is the only model checker which verifies MAS directly applying SAT-based methods, and the only one handling formulas which combine knowledge and real time. In addition, VerICS is also able to verify RTS. The results for BMC, as well as preliminary results for UMC (obtained using our modification of the SAT-solver ZChaff), seem to be quite promising, taking into account that the systems tested are not optimised w.r.t. the properties to be verified.

References

1. R. Bordini, M. Fisher, C. Pardavila, W. Visser, and M. Wooldridge. Model checking multi-agent programs with CASP. In *Proc. of the Int. Conf. on Computer Aided Verification (CAV'03)*, vol. 2725 of *LNCS*, pp. 110–113. Springer-Verlag, 2003.
2. E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
3. P. Dembiński, A. Janowska, P. Janowski, W. Penczek, A. Pólrola, M. Szreter, B. Woźna, and A. Zbrzezny. VerICS: A tool for verifying timed automata and Estelle specifications. In *Proc. of the 9th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'03)*, vol. 2619 of *LNCS*, pp. 278–283. Springer-Verlag, 2003.
4. E. Emerson and E. Clarke. Characterizing correctness properties of parallel programs using fixpoints. In *Proc. of the 7th Int. Colloq. on Automata, Languages and Programming (ICALP'80)*, vol. 85 of *LNCS*, pp. 169–181. Springer-Verlag, 1980.
5. R. Fagin, J. Y. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, 1995.
6. P. Gammie and R. van der Meyden. MCK: Model checking the logic of knowledge. In *Proc. of the 16th Int. Conf. on Computer Aided Verification (CAV'04)*, vol. 3114 of *LNCS*, pp. 479–483. Springer-Verlag, 2004.

7. ISO/IEC 9074(E), Estelle - a formal description technique based on an extended state-transition model. International Standards Organization, 1997.
8. M. Kacprzak, A. Lomuscio, T. Lasica, W. Penczek, and M. Szreter. Verifying multiagent systems via unbounded model checking. In *Proc. of the 3rd NASA Workshop on Formal Approaches to Agent-Based Systems (FAABS III)*, vol. 3228 of *LNCIS*, pp. 189–212. Springer-Verlag, 2005.
9. M. Kacprzak, A. Lomuscio, A. Niewiadomski, W. Penczek, F. Raimondi, and M. Szreter. Comparing BDD and SAT based techniques for model checking Chaum’s dining cryptographers protocol. *Fundam. Informaticae*, 72(1-2):215–234, 2006.
10. M. Kacprzak, A. Lomuscio, and W. Penczek. From bounded to unbounded model checking for temporal epistemic logic. *Fundam. Informaticae*, 63(2-3):221–240, 2004.
11. A. Lomuscio and F. Raimondi. MCMAS: A model checker for multi-agent systems. In *Proc. of the 12th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’06)*, vol. 3920 of *LNCIS*, pp. 450–454. Springer-Verlag, 2006.
12. A. Lomuscio, B. Woźna, and A. Zbrzezny. Bounded model checking real-time multi-agent systems with clock differences: Theory and implementation. In *Post-proc. of the 4th Int. Workshop on Model Checking and Artificial Intelligence (MoChArt’06)*, vol. 4428 of *LNAI*, pp. 96–112. Springer-Verlag, 2007.
13. K. L. McMillan. Applying SAT methods in unbounded symbolic model checking. In *Proc. of the 14th Int. Conf. on Computer Aided Verification (CAV’02)*, vol. 2404 of *LNCIS*, pp. 250–264. Springer-Verlag, 2002.
14. MiniSat. <http://www.cs.chalmers.se/Cs/Research/FormalMethods/MiniSat>, 2006.
15. W. Penczek and A. Lomuscio. Verifying epistemic properties of multi-agent systems via bounded model checking. In *Proc. of the 2nd Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS’03)*, pp. 209–216. ACM, 2003.
16. W. Penczek and A. Pórola. *Advances in Verification of Time Petri Nets and Timed Automata: A Temporal Logic Approach*, vol. 20 of *Studies in Computational Intelligence*. Springer-Verlag, 2006.
17. RSat. <http://reasoning.cs.ucla.edu/rsat>, 2006.
18. M. Szreter. *SAT-Based Model Checking of Distributed Systems*. PhD thesis, ICS PAS, January 2007.
19. A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.
20. M. Wooldridge, M. Fisher, M. P. Huget, and S. Parsons. Model checking multiagent systems with MABLE. In *Proc. of the 1st Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS’02)*, vol. II, pp. 952–959. ACM, 2002.
21. B. Woźna, A. Lomuscio, and W. Penczek. Bounded model checking for deontic interpreted systems. In *Proc. of the 2nd Int. Workshop on Logic and Communication in Multi-Agent Systems (LCMAS’04)*, vol. 126 of *ENTCS*, pp. 93–114. Elsevier, 2005.
22. B. Woźna, A. Lomuscio, and W. Penczek. Bounded model checking for knowledge and real time. In *Proc. of the 4th Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS’05)*, pp. 165–172. ACM, 2005.
23. A. Zbrzezny. Improvements in SAT-based reachability analysis for timed automata. *Fundam. Informaticae*, 60(1-4):417–434, 2004.
24. A. Zbrzezny. SAT-based reachability checking for timed automata with diagonal constraints. *Fundam. Informaticae*, 67(1-3):303–322, 2005.
25. L. Zhang. Zchaff. <http://www.ee.princeton.edu/~chaff/zchaff.php>, 2001.