

Checking Interval Based Properties for Reactive Systems

Pei Yu^{1*} and Xu Qiwen²

¹ International Institute for Software Technology
United Nations University, Macao SAR, P.R. China
peiyu@iist.unu.edu

² Faculty of Science and Technology
University of Macau, Macao SAR, P.R. China
qwxu@umac.mo

Abstract. A reactive system does not terminate and its behaviors are typically defined as a set of infinite sequences of states. In formal verification, a requirement is usually expressed in a logic, and when the models of the logic are also defined as infinite sequences, such as the case for LTL, the satisfaction relation is simply defined by the containment between the set of system behaviors and that of logic models. However, this satisfaction relation does not work for interval temporal logics, where the models can be considered as a set of finite sequences. In this paper, we observe that for different interval based properties, different satisfaction relations are sensible. Two classes of properties are discussed, and accordingly two satisfaction relations are defined, and they are subsequently unified by a more general definition. A tool is developed based on the Spin model checking system to verify the proposed general satisfaction relation for a decidable subset of Discrete Time Duration Calculus.

Keywords: model checking, finitary property, reactive system, interval temporal logic

1 Introduction

A reactive system does not terminate and its semantics is typically defined as the set of infinite sequences of states generated by the execution of the system. In formal verification, a requirement is usually expressed in a logic. A popular specification logic for reactive systems is LTL, and it is not by chance that a model of an LTL formula is also an infinite sequence of states. In this setting, the satisfaction relation is simply the containment between the set of system behaviors and that of logic models.

Another class of temporal logic is interval based. Although interval logics are less widely used compared to LTL, there are properties that are easier to express in such logics [14]. In almost all the interval based logics, intervals are finite ones, corresponding to finite sequences of states, which makes set containment

* On leave from Department of Computer Science and Technology, Nanjing University, P. R. China

no longer an appropriate satisfaction relation when these logics are used as the specification logic. One possible definition for this satisfaction relation is that a reactive system satisfies such a property iff all the finite prefixes of all the infinite behaviors satisfy the property, i.e. all the prefixes must be models of the requirement formula. However, while this definition makes sense in some cases, it does not in some other cases. This depends on the kind of properties. Take a mutual exclusion algorithm as an example, there are two well-known desirable properties: no two processes can be in critical sections at the same time, and a process should be allowed to enter the critical section eventually. An attempt to express these two properties in the Interval Temporal Logic (ITL) [10] or one of its variants may result in the following two formulas:

- $\Box\neg(ain \wedge bin)$: a model of it is a finite interval such that *ain* and *bin* are never true together in any of the sub-interval;
- $\Diamond ain$: a model of it is a finite interval such that *ain* is true in one state.

For $\Box\neg(ain \wedge bin)$, all the execution prefixes of a mutual exclusion algorithm should be its models, and if all the execution prefixes are models of the formula, intuitively the first requirement is satisfied. This is due to the fact that the property is a *safety property* [9], for which the above mentioned satisfaction relation is appropriate. On the other hand, for $\Diamond ain$, this satisfaction relation is obviously not appropriate - clearly we cannot expect for all the execution prefixes of a mutual exclusion algorithm, a particular process is in critical section. The second property is a *liveness property* [9].

There are some efforts to develop logics over infinite intervals [11, 16, 19]. With such logics, it is possible to use set containment as the satisfaction relation. However, logics over infinite intervals have not been widely accepted due to several reasons. Firstly, length of interval is an important concept in ITL, but for an infinite interval, this leads to an infinite number, which many people do not feel comfortable to deal with. Secondly, ITL can be decided by using finite automata when it is over finite intervals [5, 12], but obviously needs automata over infinite words when the logic is defined over infinite intervals. Even worse, effective automaton constructing techniques for LTL do not seem to apply to interval logics as observed by Wolper [17].

In this paper, we propose to use the usual interval logic over finite intervals as the specification logic, and redefine the satisfaction relation. We identify two classes of properties, and accordingly define two satisfaction relations. The two relations are unified by a more general definition, and a method to automatically check whether a reactive system satisfies properties in interval logics according to this new definition is investigated. This turns out to be checking containment between the set of (infinite) behaviors of the reactive system and the language of a Büchi automaton which is the same as the finite automaton for the interval logic formula except the final states of the finite automaton are now the accepting states of the Büchi automaton.

As a particular application of our method, we have chosen a decidable subset of Discrete Time Duration Calculus [5] as the specification logic and developed a

model checking tool, integrated with the Spin model checking system, to verify the proposed general satisfaction relation.

The rest of the paper is organized as follows. In the next section, we give the syntax and the semantics of the logic, and the construction that converts a formula to a regular language representing the models of the formula. In Section 3, two kinds of properties, finitary safety property and eventual persistence property, are defined, followed by methods to recognize them. Satisfaction relations for finitary safety property and eventual persistence property are first studied separately and then unified in Section 4. A method to model check a reactive system with the proposed satisfaction relation is shown in Section 5. In Section 6, as a case study, Peterson's mutual exclusion algorithm is model checked. The paper ends with a discussion.

2 Quantified RDC

Quantified RDC(QRDC) is an extension of the decidable subset of Duration Calculus by introducing quantifiers on state variables.

We use the following notations throughout the paper. For a finite set Σ , Σ^+ denotes the set of all non-empty finite sequences over Σ and $\Sigma^* = \Sigma^+ \cup \{\varepsilon\}$, where ε is the unique empty sequence. We use Σ^ω to denote the set of infinite sequences, i.e. ω -sequences, over Σ . $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$. For a sequence α over Σ , $|\alpha|$ denotes the length of α . Specially, $|\alpha| = \infty$ for $\alpha \in \Sigma^\omega$ and $|\varepsilon| = 0$. For two sequences $\alpha_1 \in \Sigma^*$ and $\alpha_2 \in \Sigma^\infty$, $\alpha_1 \cdot \alpha_2$ denotes the sequence obtained by concatenating α_2 to the end of α_1 ; α_1 is called a prefix of α_2 , denoted as $\alpha_1 \preceq \alpha_2$, iff $\exists \alpha_3 \in \Sigma^\infty : \alpha_1 \cdot \alpha_3 = \alpha_2$. For two sets of finite sequences P_1 and P_2 , $P_1 \cdot P_2 = \{\alpha_1 \cdot \alpha_2 \mid \alpha_1 \in P_1 \wedge \alpha_2 \in P_2\}$.

2.1 Syntax

We use p, p_1, \dots, p_n to denote state variables and $SVar$ the finite set of all the state variables. State expressions and formulas are defined as follows:

$$\begin{aligned} S &::= 0 \mid 1 \mid p \mid \neg S_1 \mid S_1 \vee S_2 \\ \phi &::= \llbracket S \rrbracket \mid \neg \phi \mid \phi_1 \vee \phi_2 \mid \phi_1 ; \phi_2 \mid \exists p. \phi \end{aligned}$$

2.2 Semantics

The interpretation for state variables is given as a function

$$I \in SVar \rightarrow (Time \rightarrow \{0, 1\})$$

When we choose $Time$ to be \mathbb{N} , we obtain QRDC in discrete time (DQRDC). The interpretation function I is extended to state expression S by induction on the structure of S :

1. $I\llbracket 0 \rrbracket(t) = 0$
2. $I\llbracket 1 \rrbracket(t) = 1$
3. $I\llbracket p \rrbracket(t) = I(p)(t)$

4. $I[\neg S](t) = 1 - I[S](t)$
5. $I[S_1 \vee S_2](t) = \max(I[S_1](t), I[S_2](t))$

Given an interpretation I , the semantics of a QRDC formula ϕ is given by a function

$$I[\phi] \in \mathbb{Intv} \rightarrow \{tt, ff\},$$

where the set of time intervals is defined as

$$\mathbb{Intv} \hat{=} \{[b, e] \mid b, e \in \mathbb{Time} \wedge b \leq e\}.$$

The function is defined as follows:

1. $I[\llbracket S \rrbracket](b, e) = tt$ iff $b < e$ and $\int_b^e I[S](t) = e - b$ for dense time or $b < e$ and $\sum_{t=b}^{e-1} I[S](t) = e - b$ for discrete time
2. $I[\neg\phi](b, e) = tt$ iff $I[\phi](b, e) = ff$
3. $I[\phi_1 \vee \phi_2](b, e) = tt$ iff $I[\phi_1](b, e) = tt$ or $I[\phi_2](b, e) = tt$
4. $I[\phi_1; \phi_2](b, e) = tt$ iff $I[\phi_1](b, m) = tt$ and $I[\phi_2](m, e) = tt$, for some $m \in [b, e] \cap \mathbb{Time}$
5. $I[\exists p. \phi](b, e) = tt$ iff for some interpretation I' , which is p -equivalent to I , $I'[\phi](b, e) = tt$.

An interpretation I' is p -equivalent to another interpretation I iff $I(p_1)(t) = I'(p_1)(t)$ for all $p_1 \neq p$ and $t \in \mathbb{Time}$.

Standard abbreviations from predicate logic, e.g. “ \wedge ”, “ \Rightarrow ” and “ \Leftrightarrow ”, are used in this paper. Moreover, the following abbreviations for QRDC formulas are frequently used:

$$\begin{array}{ll} \llbracket S \rrbracket \hat{=} \neg \llbracket 1 \rrbracket & \llbracket S \rrbracket^* \hat{=} \llbracket S \rrbracket \vee \llbracket \rrbracket \\ false \hat{=} \llbracket 0 \rrbracket & true \hat{=} \neg false \\ \diamond\phi \hat{=} true; \phi; true & \square\phi \hat{=} \neg\diamond\neg\phi \\ \square_p\phi \hat{=} \neg(\neg\phi; true) & l = 0 \hat{=} \neg \llbracket 1 \rrbracket \end{array}$$

For DQRDC, we also use the following abbreviations ($k \in \mathbb{N}$):

$$\begin{array}{ll} l = 1 \hat{=} \llbracket 1 \rrbracket \wedge \neg(\llbracket 1 \rrbracket; \llbracket 1 \rrbracket) & fS = 1 \hat{=} (fS = 0); (\llbracket S \rrbracket \wedge l = 1); (fS = 0) \\ fS = 0 \hat{=} \llbracket \neg S \rrbracket \vee l = 0 & fS = k + 1 \hat{=} (fS = k); (fS = 1) \quad (k \geq 1) \\ fS \geq k \hat{=} (fS = k); true & fS > k \hat{=} fS \geq k + 1 \\ fS \leq k \hat{=} \neg(fS > k) & fS < k \hat{=} fS \leq k - 1 \end{array}$$

where fS denotes the accumulated time when S evaluates 1. The length l of current interval can be encoded as $f1$.

2.3 Segments as Models

We call each $obs \in Obs = 2^{SVar}$ an observation. For a state expression S , $N(S) \subseteq Obs$ denotes the set of observations where S is evaluated to 1:

$$N(S) \hat{=} \{obs \mid \bigwedge_{p_1 \in obs} p_1 \wedge \bigwedge_{p_2 \in (SVar - obs)} \neg p_2 \Rightarrow S\}.$$

For an interpretation I and $t \in \mathbb{Time}$, $I(t)$ is identified with the following observation

$$I(t) \hat{=} \{p \mid p \in SVar \wedge I(p)(t) = 1\}$$

In DQRDC, given an interpretation I and an interval $[b, e]$, we use the pair $(I, [b, e])$, called a segment, to denote the finite sequence $I(b), I(b+1), \dots, I(e-1)$ of observations. A segment $(I, [b, e])$ satisfies a formula ϕ , denoted as $I, [b, e] \models \phi$, iff $I[\![\phi]\!](b, e) = tt$. For a formula ϕ , we call each segment $(I, [b, e])$ satisfying $I, [b, e] \models \phi$ a model of ϕ . The set of models of ϕ , denoted as $\mathcal{M}[\![\phi]\!]$, can be constructed [5, 12] as the following regular language over the alphabet Obs :

$$\begin{aligned}
\mathcal{M}[\![S]\!] &= (N(S))^+ && \text{(positive closure)} \\
\mathcal{M}[\![\phi_1 \vee \phi_2]\!] &= \mathcal{M}[\![\phi_1]\!] \cup \mathcal{M}[\![\phi_2]\!] && \text{(union)} \\
\mathcal{M}[\![\neg\phi]\!] &= Obs^* - \mathcal{M}[\![\phi]\!] && \text{(complementation)} \\
\mathcal{M}[\![\phi_1; \phi_2]\!] &= \mathcal{M}[\![\phi_1]\!] \cdot \mathcal{M}[\![\phi_2]\!] && \text{(concatenation)} \\
\mathcal{M}[\![\exists p. \phi]\!] &= Equival_p(\mathcal{M}[\![\phi]\!]) && \text{(equivalence closure)}
\end{aligned}$$

where $Equival_p(\mathcal{M}[\![\phi]\!])$ is defined as follows.

Definition 1. Given two finite sequences $\alpha = s_0, s_1, \dots, s_n$ and $\alpha' = s'_0, s'_1, \dots, s'_n$ of observations and a state variable p , α and α' are p -equivalent iff $s_i \setminus \{p\} = s'_i \setminus \{p\}$ for $0 \leq i \leq n$.

Given a state variable p , the equivalence closure of a set of observation sequences \mathcal{M} , $Equival_p(\mathcal{M})$ is defined to be the set

$$\{\alpha \mid \text{if there exists } \beta \in \mathcal{M}, \text{ such that } \alpha \text{ and } \beta \text{ are } p\text{-equivalent}\}$$

We can easily prove the following lemma by induction on the structure of ϕ .

Lemma 2. In DQRDC, for a formula ϕ and a segment $(I, [b, e])$, $(I, [b, e]) \in \mathcal{M}[\![\phi]\!]$ iff $I, [b, e] \models \phi$.

In DQRDC, a formula ϕ is valid iff $\mathcal{M}[\![\phi]\!] = Obs^*$, and it is satisfiable iff $\mathcal{M}[\![\phi]\!] \neq \emptyset$.

3 Finitary Safety Property and Eventual Persistence Property

Our work on classification of properties follows that of [9, 3]. However, in our setting, a property is a set of finite sequences of states, and therefore some modification is necessary. To indicate whether properties are formed by finite or infinite sequences of states, we call them finitary properties or infinitary properties.

3.1 Finitary Safety Property

A safety property stipulates that “something (bad) never happens”. In [3], it is considered that if a “bad” thing happens in an infinite sequence, then it must do so after some finite prefix and must be irremediable. We follow this view, but since our property is over finite sequences, we call our property *finitary safety property* and it is formally defined as follows:

Definition 3 (finitary safety property). For a finite set Σ and a property $P \subseteq \Sigma^*$, P is a finitary safety property, or an FS property, iff:

$$\forall \alpha \in \Sigma^* : (\alpha \notin P \Leftrightarrow \exists \beta \in \Sigma^* : (\beta \preceq \alpha \wedge \forall \gamma \in \Sigma^* : (\beta \preceq \gamma \Rightarrow \gamma \notin P))).$$

By taking negation on both sides, we have

$$\forall \alpha \in \Sigma^* : (\alpha \in P \Leftrightarrow \forall \beta \in \Sigma^* : (\beta \preceq \alpha \Rightarrow \exists \gamma \in \Sigma^* : (\beta \preceq \gamma \wedge \gamma \in P))).$$

This can be simplified, and we have the following theorem which says that a property is an FS property iff it is prefix-closed.

Theorem 4. A non-empty property $P \subseteq \Sigma^*$ is an FS property iff

$$\forall \alpha \in \Sigma^* : (\alpha \in P \Leftrightarrow \forall \beta \in \Sigma^* : (\beta \preceq \alpha \Rightarrow \beta \in P)),$$

i.e. $\text{Pref}(P) = P$, where $\text{Pref}(P) = \{\beta \mid \beta \in \Sigma^* \wedge \exists \alpha \in P : \beta \preceq \alpha\}$.

3.2 Eventual Persistence Property

An infinitary property is a liveness property [3] iff all finite executions can be extended to satisfy the property. Intuitively, this is to say a “good” thing always has the chance to take place no matter what has happened. Although in the general case, the “good thing” may take infinite number of actions to ensure, for example, something that happens infinitely often, some “good thing” can and in fact is only meaningful to occur after a finite number of steps. This is often called *eventuality*. Moreover, we are particularly interested in a class of properties such that if something (“good”) has happened, nothing later can undo it. This is known as *persistence*. Formally, eventuality and persistence properties are defined as:

Definition 5 (Eventuality and Persistence). For a finite set Σ and a property $P \subseteq \Sigma^*$ ($P \neq \emptyset$),

- P is an eventuality property iff $\forall \alpha \in \Sigma^* : \exists \beta \in \Sigma^* : (\alpha \preceq \beta \wedge \beta \in P)$;
- P is a persistence property iff $\forall \alpha \in P : \forall \beta \in \Sigma^* : (\alpha \preceq \beta \Rightarrow \beta \in P)$.

P is called an eventual persistence property, or an EP property, iff P is both an eventuality property and a persistence property.

The following theorem follows directly from the definition.

Theorem 6. A non-empty property $P \subseteq \Sigma^*$ is an EP property iff $\text{Pref}(P) = \Sigma^*$ and $P \cdot \Sigma^* = P$.

3.3 Recognizing FS and EP Properties

A finitary property can be specified by a finite automaton whose language is the set of the finite sequences satisfying the property. In the following, similar to what has been done in [2] w.r.t. Büchi automata expressing safety properties and liveness properties, we present the rules for recognizing FS and EP properties from automata theory point of view.

A finite automaton FA is reduced if, for any state q , there is a path from an initial state to q and there is a path from q to a final state. From a finite automaton, a reduced FA accepting the same language can always be obtained by deleting those states either not reachable from any initial state or from which no final state can be reached. For a reduced FA , we define its closure $cl(FA)$ as the result automaton by setting all of its states as final states. From the construction of $cl(FA)$, we can easily get

Lemma 7. *For a reduced FA , $L(cl(FA)) = Pref(L(FA))$.*

We have from Theorem 4 and Lemma 7 the following Theorem for FS properties:

Theorem 8. *A reduced finite automaton FA specifies an FS property iff $L(FA) = L(cl(FA))$.*

For EP properties, we have the following similar theorem, which follows from Theorem 6 and Lemma 7:

Theorem 9. *A finite automaton FA specifies an EP property iff $L(cl(FA)) = \Sigma^*$ and $L(FA) \cdot \Sigma^* = L(FA)$.*

From Theorems 8 and 9, algorithms for recognizing FS and EP properties have been developed [18].

3.4 FS Formulas and EP Formulas in DQRDC

In DQRDC, each formula ϕ defines a finitary property $\mathcal{M}[\phi]$. ϕ is called an FS/EP formula iff $\mathcal{M}[\phi]$ is an FS/EP property. For any DQRDC formula, we can always semantically identify whether it is an FS/EP formula or not by examining the automaton accepting $\mathcal{M}[\phi]$. The following two theorems provide syntactical conditions which are sufficient (but not necessary) to judge these two kinds of formulas. Due to limitation of space, the detailed proofs are omitted, and can be found in [18].

Theorem 10 (FS formula). *Every DQRDC formula of the form $\llbracket S \rrbracket^*$, $\int S \leq k$ ($k \in \mathbb{N}$), $\Box_p \phi$ or $\Box \phi$ is an FS formula and if ϕ_1, ϕ_2 are FS formulas, then so are $\phi_1 \wedge \phi_2$, $\phi_1 \vee \phi_2$, $\phi_1; \phi_2$ and $\exists p. \phi_1$.*

Theorem 11 (EP formula). *Every DQRDC formula of the form $\int S \geq k$ ($S \neq 0$ and $k \in \mathbb{N}$) and $\Diamond \phi$ is an EP formula and if ϕ_1, ϕ_2 are EP formulas, then so are $\phi_1 \wedge \phi_2$, $\phi_1 \vee \phi_2$, $\phi_1; \phi_2$ and $\exists p. \phi_1$.*

4 Satisfaction Relations

A finite sequence α satisfies a finitary property P , denoted by $\alpha \models P$, iff $\alpha \in P$. For an ω -sequence, we first look at the satisfaction relations for FS and EP properties separately.

For an FS property, an ω -sequence satisfies the property iff all the finite prefixes of the sequence satisfy it.

Definition 12. For an ω -sequence σ and an FS property P all defined over Σ , σ satisfies P , denoted as $\sigma \models_i P$, iff $\forall \alpha \in \Sigma^* : (\alpha \preceq \sigma \Rightarrow \alpha \in P)$.

Intuitively, an ω -sequence satisfies an EP property iff it has a finite prefix where the “good” thing has happened.

Definition 13. For an ω -sequence σ and an EP property P all defined over Σ , σ satisfies P , denoted as $\sigma \models_i P$, iff $\exists \alpha \in \Sigma^* : (\alpha \preceq \sigma \wedge \alpha \in P)$.

The two satisfaction relations can be unified by a more general one.

Definition 14 (satisfaction of finitary property). For an ω -sequence σ and a finitary property P , both defined over Σ , σ satisfies P , denoted as $\sigma \models_i P$, iff $\exists \gamma \in \Sigma^* : (\gamma \preceq \sigma \wedge \forall \alpha \in \Sigma^* : (\gamma \preceq \alpha \preceq \sigma \Rightarrow \alpha \models P))$.

Intuitively, this says that an ω -sequence satisfies a finitary property if all except a finite number of the finite prefixes satisfy the property.

Lemma 15. Definition 14 is equivalent to Definition 12 for FS properties and Definition 13 for EP properties.

Proof. For an FS property P , Theorem 4 says that P is prefix-closed, and it follows that for any ω -sequence σ over Σ
 $\forall \alpha \in \Sigma^* : (\alpha \preceq \sigma \Rightarrow \alpha \in P)$ iff $\exists \gamma \in \Sigma^* : (\gamma \preceq \sigma \wedge \forall \alpha \in \Sigma^* : (\gamma \preceq \alpha \preceq \sigma \Rightarrow \alpha \in P))$.
 Namely, Definitions 12 and 14 are equivalent for FS properties.

For an EP property P , Definition 5 indicates that if a prefix satisfies P , then all its extensions satisfy P too. Therefore, for any ω -sequence σ over Σ
 $\exists \alpha \in \Sigma^* : (\alpha \preceq \sigma \wedge \alpha \in P)$ iff $\exists \gamma \in \Sigma^* : (\gamma \preceq \sigma \wedge \forall \alpha \in \Sigma^* : (\gamma \preceq \alpha \preceq \sigma \Rightarrow \alpha \in P))$.
 Subsequently, Definitions 13 and 14 are equivalent for EP properties. \square

5 Model Checking

Model checking [4] is a verification technique that involves constructing a model of a system and checking if a desired property holds in that model. In this section, we study model checking finitary properties specified in DQRDC.

For a system S , $\mathcal{H}[S]$ denotes the set of all behaviors, i.e., ω -sequences of system states, of S , and S satisfies a DQRDC formula ϕ expressing finitary property iff $\forall \sigma \in \mathcal{H}[S] : \sigma \models_i \phi$, i.e.

$$\forall \sigma \in \mathcal{H}[S] : (\exists \gamma \in Obs^* : (\gamma \preceq \sigma \wedge \forall \alpha \in Obs^* : (\gamma \preceq \alpha \preceq \sigma \Rightarrow \alpha \models \phi))).$$

Subsequently, S does not satisfy ϕ iff

$$\exists \sigma \in \mathcal{H}[S] : (\forall \gamma \in Obs^* : (\gamma \preceq \sigma \Rightarrow \exists \alpha \in Obs^* : (\gamma \preceq \alpha \preceq \sigma \wedge \alpha \not\models \phi))).$$

We can construct a finite automaton $FA(\neg\phi)$, which accepts precisely all the finite sequences not satisfying ϕ , thus the above condition can be rewritten as $\exists \sigma \in \mathcal{H}[S] : (\forall \gamma \in Obs^* : (\gamma \preceq \sigma \Rightarrow \exists \alpha \in Obs^* : (\gamma \preceq \alpha \preceq \sigma \wedge \alpha \in L(FA(\neg\phi)))))$. When $FA(\neg\phi)$ is deterministic, referred to as $DFA(\neg\phi)$, we shall prove this condition is equivalent to

$$\exists \sigma \in \mathcal{H}[S] : \sigma \in L(DBA(\neg\phi)),$$

where $DBA(\neg\phi)$ is the deterministic Büchi automaton [1] obtained by taking final states of $DFA(\neg\phi)$ as accepting states.

Therefore, to determine whether S satisfies ϕ , it suffices to check whether $\mathcal{H}[S] \cap L(DBA(\neg\phi)) = \emptyset$. If the intersection is empty, S satisfies ϕ , and if not, S does not satisfy ϕ .

Lemma 16. *For an ω -sequence σ and a deterministic finite automaton DFA , let DBA be the deterministic Büchi automaton obtained by taking all the final states of DFA as accepting states, then*

$$\forall \alpha \in Obs^* : \alpha \preceq \sigma : (\exists \gamma \in Obs^* : \alpha \preceq \gamma \preceq \sigma : \alpha \in L(DFA)) \text{ iff } \sigma \in L(DBA).$$

Proof. The if case is obvious and we only prove the only if case. From the assumption, there exists an infinite series of finite sequences $\alpha_1, \alpha_2, \dots$, with $\alpha_1 \preceq \alpha_2 \preceq \dots \preceq \sigma$, $\alpha_i \neq \alpha_{i+1}$ and $\alpha_i \in L(DFA)$ for all $i \geq 1$. Obviously, α_1 ends in a final state. Because the automaton is deterministic, α_2 will first go through the same states as α_1 , in particular, the first final state, and then ends in a (possibly different) final state. Therefore, α_2 goes through the final states which are the accepting states of the Büchi automaton at least twice. By induction, we know that α_i goes through the accepting states of the Büchi automaton at least i times. Since there are infinite many such α_i 's, some accepting states are gone through infinite number of times by σ , and therefore $\sigma \in L(DBA)$. \square

Note in the above proof, the assumption that the automaton is deterministic is important. However, this is not necessary for checking FS properties. If ϕ is an FS property, in $FSA(\neg\phi)$, any input symbol will lead a final state to (another) final state. Informally this is because for an FS property, if something “bad” has happened in a behavior, any extension of it is still a bad behavior. Therefore, if a prefix of σ has reached a final state, any further input will lead to a final state, and subsequently σ is accepted by the corresponding Büchi automaton. Since determinization can be expensive, it is in general more efficient to use non-deterministic automata for model checking FS properties.

Spin [7] is a widely distributed software system for on-the-fly model checking. In Spin, the system is modelled using Promela, its input language, as a set of interacting processes. The negation of the desired property is encoded in a temporal claim, which also takes the form of a Promela program. Each process is translated into an automaton, and the global behavior of the system is obtained by computing the asynchronous product of these automata. The temporal claim is translated to a Büchi automaton. The synchronous product of this automaton and the automaton representing the global behavior of the system is computed. If the language accepted by the resulting Büchi automaton is empty, no behavior of the system satisfies the negation of the desired property, i.e. the system satisfies the property. Otherwise, the language contains some system behaviors violating the desired property.

We have implemented the algorithm which takes a DQRDC formula as input and generates the temporal claim in Promela. This allows us to model check properties studied in this paper when they are expressed as DQRDC formulas using Spin. Our implementation makes use of the C++ package for finite state

machine, Grail+ [8]. The compiler construction program Parser Generator from Bumble-Bee Software Ltd. is also used.

6 Case Study

In this section, we study model checking of the well-known Peterson’s mutual exclusion algorithm. In Fig. 1, the algorithm for the case of two processes is modelled in Promela.

```

1  #define true 1
2  #define false 0
3  #define aturn 1
4  #define bturn 0
5
6  bool areq, breq, turn;
7  bool ain;
8  bool bin;
9
10 proctype a()
11 {
12   do
13   :: atomic{ areq=true;
14              turn=bturn; }
15   (breq==false || turn==aturn);
16   ain=true;
17   /* critical section cs */
18   /* assert(bin==false); */
19   ain=false;
20   areq=false;
21   od
22 }
23
24 proctype b()
25 {
26   do
27   :: atomic{ breq=true;
28              turn=aturn; }
29   (areq==false || turn==bturn);
30   bin=true;
31   /* critical section */
32   /* assert(ain==false); */
33   bin=false;
34   breq=false;
35   od
36 }
37
38 init
39 {
40   run a(); run b();
41 }

```

Fig. 1. Peterson’s mutual exclusion algorithm

The algorithm uses three global variables `turn`, `areq` and `breq`. The statements on lines 14 and 26, which can only be executed if the boolean equations in them evaluate to true, synchronize the processes. If there is only one process that applies to enter the the critical section, the entry is immediate. Variable `turn` is used to decide which process to enter the critical section when both processes have requested. Variable `ain(bin)` is used here to facilitate property specification and it is set `true` iff process `a(b)` is in the critical section.

The essential property of a mutual exclusion algorithm is naturally mutual exclusion, which can be easily specified in DQRDC as

$$Req_1 \hat{=} \square(l > 1 \Rightarrow \llbracket \neg(ain \wedge bin) \rrbracket)$$

The generated corresponding Büchi automaton in Promela is as shown in Fig. 2.

The mutual exclusion property is simple and can be expressed using assertions (as shown on lines 17 and 29) or a simple monitor process. Next we consider a more involved property,

if one process has requested to enter its critical section when the other process is in the critical section, then the first process must have entered the critical section before the second one enters the critical section for another time.

```

1 never {
2 T0_init:
3   if
4   :: goto T1
5   fi;
6 T1:
7   if
8   :: ((true))->goto T1
9   :: ((!ain || !bin))->goto T2
10  :: ((ain && bin))->goto accept_5
11  fi;
12 T2:
13  if
14  :: ((!ain || !bin))->goto T2
15  :: ((ain && bin))->goto accept_5
16  fi;
17 accept_5:
18  if
19  :: ((true))->goto accept_5
20  fi;
21 }

```

Fig. 2. Temporal claim for Req_1 .

This property can be expressed as follows w.l.o.g.:

$$Req_2 \hat{=} \Box(\llbracket areq \wedge bin \rrbracket; \llbracket \neg bin \rrbracket; \llbracket bin \rrbracket \Rightarrow \Diamond \llbracket ain \rrbracket)$$

The generated Büchi automaton in Promela is shown in Fig. 3.

```

1 never {
2 T0_init:
3   if
4   :: goto T0
5   fi;
6 T0:
7   if
8   :: ((true)) -> goto T0
9   :: ((areq && bin && !ain))->goto T2
10  fi;
11 T2:
12  if
13  :: ((!bin && !ain)) -> goto T6
14  :: ((areq && bin && !ain))->goto T2
15  fi;
16 T6:
17  if
18  :: ((!bin && !ain)) -> goto T6
19  :: ((bin && !ain)) -> goto accept_14
20  fi;
21 accept_14:
22  if
23  :: ((true)) -> goto accept_19
24  :: ((bin && !ain)) -> goto accept_14
25  fi;
26 accept_19:
27  if
28  :: ((true)) -> goto accept_19
29  fi;
30 }

```

Fig. 3. Temporal claim for Req_2 .

Taking the temporal claim and the system model as input, the Spin model checker shows that the properties are satisfied.

7 Discussion

In this paper, we have studied the automatic verification of reactive systems against interval logic specifications. Although various interval logics have been developed for some time, there is little work with similar goal, which is somewhat unusual in view of extensive research regarding LTL. In fact, there have been no clear definitions about what it means for a reactive system to satisfy properties in interval logics. We have considered two classes of common properties and for them defined satisfaction relations that have natural meanings. The two satisfaction relations are unified by a more general one for which model checking support has been provided. We are not clear if the general relation is meaningful for properties outside the two classes and their simple combinations.

Our work is useful to people who prefer to use interval logics. In fact, both LTL and interval logics can be used together in specifying properties of a system.

For example, in Peterson's mutual exclusion algorithm case study, one can use the interval logic to specify Req_2 as we have done, since this property may not be easy to express in LTL, but use LTL to specify other properties for which LTL is better or just as good.

Although the algorithm converting the interval logic formulas to finite automata is simple and well-known, there have been few implementations. The early BDD-based implementation by Skakkebak and Sestoft is integrated into a proof assistant [15] developed on an old version of PVS which has not been supported for years. The more recent implementation by Pandya [12] uses the Mona [6] system by mapping the interval logic to the language of Mona, a monadic second-order logic. Although the implementation of Mona is highly efficient, its sophistication also makes it difficult to handle. What we really need is a package on finite automata, and with Grail+, we have a better control of the implementation. For example, we can control whether to determinize the automaton, which is an expensive operation.

Pandya has also studied model checking reactive systems with interval logics. In his approach, the interval logic does not stand alone and is combined into other logics, for example, into CTL [13]. Pandya has apparently considered combining the interval logic with LTL, and using Spin to check properties expressed in the new logic. The resulting logics are substantially different from existing ones and considerable work may be needed before they are accepted.

Acknowledgements: This research is partly supported by University Macau Research Grant No. RG039/02-038/XQW/FST.

References

1. J.R. Buchi. On a Decision Method in Restricted Second Order Arithmetic. In Nagel et al., editor, Logic, Methodology and Philosophy of Science. Stanford Univ. Press, 1960.
2. Bowen Alpern and Fred B. Schneider. Recognizing Safety and Liveness. TR 86-727, 1986.
3. Bowen Alpern and Fred B. Schneider. Defining Liveness. Information Processing Letters, 21:181–185, 1985.
4. Edmund M. Clarke, Orna Grumberg and Doron A. Peled. Model Checking. The MIT Press, 1999.
5. Michael R. Hansen and Zhou Chaochen. Duration Calculus: Logical Foundations. Formal Aspects of Computing, 9:283-330, 1997.
6. J.G. Henriksen, J. Jensen, M. Jorgensen, N. Klarlund, B. Paige, T. Rauhe, and A. Sandholm. Mona: Monadic Second-order Logic in practice. In Proc. of TACAS'95, LNCS 1019, Springer-Verlag, 1996.
7. G. Holzmann. The SPIN Model Checker. IEEE Trans. on Software Engineering, 23:279-295, 1997
8. The Grail+ Project. Department of Computer Science, University of Western Ontario, Canada. <http://www.csd.uwo.ca/research/grail/>
9. Leslie Lamport. Proving the correctness of multiprocess programs. IEEE Transactions on Software Engineering, SE-3(2):125-143, March 1977.

10. B. Moszkowski. A Temporal Logic for Multilevel Reasoning about Hardware. *IEEE Computer*, 18(2):10-19, 1985.
11. B. Moszkowski. Compositional reasoning about projected and infinite time. *Proc. of the First IEEE International Conference on Engineering of Complex Computer Systems (ICECCS'95)*, IEEE Computer Society Press, p238-245, 1995.
12. P.K. Pandya. Specifying and Deciding Quantified Discrete-Time Duration Calculus Formulae using DCVALID. In *Proc. of Workshop on Real-time Tools*, Aalborg, Denmark, 2001. TCS00-PKP-1, Tata Institute of Fundamental Research, Mumbai, India, 2000.
13. P.K. Pandya. Model checking CTL[DC]. In *Proc. of TACAS 2001*, Genova, Italy. LNCS 2031, Springer-Verlag, 2001.
14. Y. S. Ramakrishna, P. M. Melliar-Smith, Louise E. Moser, Laura K. Dillon, G. Kuttys. Interval Logics and Their Decision Procedures, Part I: An Interval Logic. *Theoretical Computer Science* 166(1&2): 1-47, Elsevier (1996).
15. J.U. Skakkebak. A Verification Assistant for Real-time Logic. PhD. Thesis, Department of Computer Science, Technical University of Denmark, 1994.
16. Wang Hanpin and Xu Qiwen. Completeness of temporal logics over infinite intervals. Technical Report 158, UNU/IIST, Macau. 1999. Accepted by *Applied Discrete Mathematics*, Elsevier.
17. P. Wolper. Constructing automata from temporal logic formula: a tutorial. *First EEF/Euro Summer School on Trends in Computer Science*, LNCS 2090, pp. 261-277, Springer-Verlag, 2001.
18. Pei Yu and Xu Qiwen. Checking Interval Based Properties for Reactive Systems. Technical Report 283, UNU/IIST, Macau. 2003.
19. Zhou Chaochen, Dang Van Hung and Li Xiaoshan. A Duration Calculus with Infinite Intervals, LNCS 965, pp. 16-41, February, 1995.