

From Offline toward Real-Time: A Hybrid Systems Model Checking and CPS Co-Design Approach for Medical Device Plug-and-Play Collaborations

Tao Li, *Student Member, IEEE*, Feng Tan, Qixin Wang, *Member, IEEE*, Lei Bu, *Member, IEEE*, Jian-nong Cao, *Senior Member, IEEE*, and Xue Liu, *Member, IEEE*,

Abstract—Hybrid systems model checking is a great success in guaranteeing the safety of computerized control *cyber-physical systems* (CPS). However, when applying hybrid systems model checking to *Medical Device Plug-and-Play* (MDPnP) CPS, we encounter two challenges due to the complexity of human body: i) there are no good offline differential equation based models for many human body parameters; ii) the complexity of human body can result in many variables, complicating the system model. In an attempt to address the challenges, we propose to alter the traditional approach of *offline* hybrid systems model checking of *time-unbounded* (i.e., *infinite-horizon*, a.k.a., long-run) future behavior to *online* hybrid systems model checking of *time-bounded* (i.e., *finite-horizon*, a.k.a., short-run) future behavior. According to this proposal, online model checking runs as a real-time task to prevent faults. To meet the real-time requirements, certain design patterns must be followed, which brings up the *co-design* issue. We propose two sets of system co-design patterns for hard real-time and soft real-time respectively. To evaluate our proposals, a case study on laser tracheotomy MDPnP is carried out. The study shows the necessity of online model checking. Furthermore, test results based on real-world human subject trace show the feasibility and effectiveness of our proposed co-design.



1 INTRODUCTION

THANKS to the rapid development of embedded systems technology, we now have thousands of kinds of embedded medical devices. So far, these devices are mainly designed for isolated use. However, people envision that by coordinating these devices, we can significantly increase medical treatment safety, capability, and efficiency. This vision led to the launch of the *Medical Device Plug-and-Play* (MDPnP) [1] effort, which aims to enable the safe composition and collaboration of disparate embedded devices in medical contexts. An MDPnP system is a typical *Cyber-Physical System* (CPS) [2]. On the one hand, it involves cyber-world discrete computer logic of various embedded medical devices. On the other hand, it involves physical-world patient-in-the-loop, which is a continuous complex biochemical system.

The top concern of any MDPnP system is safety. In the cyber-world, for a safety-critical system, people often carry out model checking [3] *before* the system is put online. In such case, model checking builds an *offline* model of the system, and checks the system's possible behaviors in the *time-unbounded future* (i.e., *infinite-horizon*). Only after passing model checking may the system be allowed to run.

This practice is a great success. For CPS verification, the state-of-the-art model checking tools are the *hybrid systems* model checking tools [4][5], which integrate the discrete automata models with the continuous differential equation (and other control theory) models. Today, hybrid systems model checking can already analyze many computerized control systems, i.e., control CPS.

The success of hybrid model checking in control CPS inspires the interest to apply it in MDPnP CPS. However, this faces a major challenge: in most MDPnP CPS, there are *no good offline* models to describe the complex biochemical system of the patient [6]. Even if some vital signs can be modeled offline, the models may not (with some exceptions [7]) fit into existing hybrid systems model checking tools, which mainly use linear differential equations to describe the physical world.

To deal with the above challenges, we propose to alter the traditional practice of *offline* model checking of hybrid system's behavior in the *infinite-horizon*. Instead, we carry out periodical *online* model checking. In every period, we only model check the hybrid system's behavior in the next (few) period(s); i.e., we only model check the hybrid system's behavior in *time-bounded future* (i.e., *finite-horizon*).

The merits of the proposed approach are as follows. First, though many human body parameters are hard to model offline, their online behaviors in finite-horizon are quite predictable. For example, after injecting 1ml of morphine, it is hard to accurately predict the blood oxygen level curve in the next 40 minutes, as it depends on too many factors, even including the patient's emotion [8][9]. However, it is easy to

- T. Li, F. Tan, Q. Wang, and J. Cao are with Dept. of Computing, The Hong Kong Polytechnic Univ., Hong Kong S. A. R. Email: csqwang@comp.polyu.edu.hk
- L. Bu is with State Key Lab for Novel Software Technology, Dept. of Computer Sci. and Tech., Nanjing Univ., Nanjing, 210093, P. R. China. Email: bulei@nju.edu.cn
- X. Liu is with School of Computer Science, McGill Univ., Montreal, Canada. Email: xue.liu@mcgill.ca

predict the blood oxygen level curve in the next 4 seconds: it cannot plunge from 100% to 10%, nor show a saw-toothed wave form; instead, it has to be smooth, which can be effectively described with existing tools, such as linear regression. Also, within short finite-horizon, we can approximate many variables as constants, and/or approximate nonlinear behaviors as linear behaviors. This would further simplify our model and computation.

The proposed approach can be formalized as follows. Given an MDPnP system \mathcal{S} , we periodically sample the observable state parameters every T seconds. At time instance kT ($k = 0, 1, 2, \dots$), we build a hybrid system model (i.e., the “online model”) of \mathcal{S} with the observed numerical values of state parameters, and verify its safety in the time interval $[kT, (k+1)T]$. We hence call T the *finite-horizon* of our online model checking. If the online model is proven safe, the system can run for another T seconds. Otherwise, the system immediately switches to an application dependant fall-back plan.

Such model checking must finish within bounded and short time, i.e. *real-time*, to allow decision making (on whether to run the system for another T seconds or switch to fall-back plan) *before* any fault happens. To support real-time, the MDPnP CPS design must follow certain patterns, which brings up the issue of hybrid systems model checking and CPS *co-design*.

In the rest of the paper, we discuss our proposed co-design approach through the context of laser tracheotomy, a representative MDPnP application [7][10]. Section 2 introduces the background on hybrid systems model checking; Section 3 proposes our online hybrid systems modeling approach; Section 4 proposes the corresponding system design patterns; Section 5 evaluates our approach; Section 6 further examines our proposal under relaxed assumptions; Section 7 discusses related work; and Section 8 concludes the paper.

This paper is based on our previous conference paper published in [11], which is in turn based on our workshop paper published in [12][13]. Compared to these previous versions, this paper mainly added Section 4.1, Theorem 3, Section 6.2, and the supplementary file.

2 BACKGROUND

Hybrid systems model checking is first proposed by Alur, Henzinger, et al. [14][15][16] and has since evolved into a family of state-of-the-art tools in CPS. The main idea is to combine the discrete automata models of computer logic with continuous differential equation models of control systems, which leads to the modeling tool of *hybrid automata*.

2.1 Syntax

Following [15]’s conventions on symbols, a hybrid automaton A is syntactically a tuple of $A = (\vec{x}, \vec{x}^0, V, v^0, inv, dif, E, act, L, syn)$, where

\vec{x} is a vector of n data variables $\vec{x} = (x_1, x_2, \dots, x_n)$. \vec{x} is regarded as a function of time, and we use $\dot{\vec{x}}$ to denote the first order derivative of \vec{x} . We also use $\vec{x}' = (x'_1, x'_2, \dots, x'_n)$ to denote the new values of \vec{x} after an event (see the definitions

for E and act). A specific evaluation of \vec{x} , denoted as $\vec{s} = (s_1, s_2, \dots, s_n) \in \mathbb{R}^n$ is called a *data state* of A . In addition, Boolean values of **true** and **false** can be denoted with real number 1 and 0 respectively; hence a data variable can also serve as a Boolean variable.

\vec{x}^0 is the initial data state.

V is a set of *locations*, a.k.a., *control locations*, where different control laws apply. Each location corresponds to a vertex in the graphical representation of hybrid automaton A . A *state* of hybrid automaton A is denoted as (v, \vec{s}) , where $v \in V$ and $\vec{s} \in \mathbb{R}^n$ is a data state.

v^0 is the initial location.

inv is the *location invariants*, a function that assigns each location $v \in V$ a set of inequalities over data variables \vec{x} . That is, when in location v , the value of \vec{x} must satisfy $inv(v)$.

dif is the *continuous activities*, a function that assigns each location $v \in V$ a set of inequalities over $\dot{\vec{x}}$ and \vec{x} . That is, when in location v , the values of $\dot{\vec{x}}$ and \vec{x} must satisfy $dif(v)$.

E is the set of *events*, a.k.a. *transitions*: edges between locations. Formally, $E \subseteq V \times V$. For an event $e = (v, v') \in E$, v is the *source location* and v' is the *target location*.

act is the *discrete actions*, a function assigns to each event $e = (v, v') \in E$ a set of inequalities over \vec{x} and \vec{x}' , where $\vec{x}' = (x'_1, x'_2, \dots, x'_n)$ refers to the new value of \vec{x} after event e . The event $e = (v, v')$ is enabled only when the value of \vec{x} in v satisfies $act(e)$, and the new value of \vec{x}' after the event is chosen nondeterministically such that $act(e)$ is satisfied. For example, suppose $\vec{x} = (x_1)$, then for $act(e) = (x_1 \leq 3 \wedge x'_1 \leq 5 \wedge x'_1 \geq 5)$, event e is only enabled when $x_1 \leq 3$; and after the event, x_1 is assigned the new value of 5. Like this example, if \vec{x} and \vec{x}' do not mix in any inequalities in $act(e)$, and \vec{x}' has a deterministic value \vec{s}' , then we can call the subset of inequalities involving only \vec{x} to be the *guard* of event e , and event e *updates* \vec{x} to \vec{s}' , denoted as $\vec{x} := \vec{s}'$.

L is a set of *synchronization labels*.

syn is the *synchronization function* that assigns each event $e \in E$ an $l \in L$. L and syn are for composition of multiple hybrid automata. Suppose we have two hybrid automata $A_1 = (\vec{x}_1, \vec{x}_1^0, V_1, v_1^0, inv_1, dif_1, E_1, act_1, L_1, syn_1)$ and $A_2 = (\vec{x}_2, \vec{x}_2^0, V_2, v_2^0, inv_2, dif_2, E_2, act_2, L_2, syn_2)$, if $e_1 \in E_1$, $e_2 \in E_2$ and $syn_1(e_1) = syn_2(e_2)$, then event e_1 and e_2 must always take place together.

Furthermore, when inv , dif , and act only involve linear inequalities, and dif does not involve \vec{x} , hybrid automaton A is called *linear hybrid automaton* (LHA)[14].

Reference [15] also describes how to combine several hybrid automata into one hybrid automaton. Particularly, the location set of the combined hybrid automaton $V_{comb} = V_1 \times V_2 \times \dots \times V_n$, where V_i ($i = 1, \dots, n$) is the location set of the i th component hybrid automaton; and “ \times ” is Cartesian product. For $v \in V_{comb}$, we use $v|_i$ to denote the projection of v on V_i .

2.2 Semantics

This paper adopts the semantic concepts and the corresponding symbol definitions of [15]. Due to page limit, interested

readers shall refer to [15] for these definitions. Of particular importance are the concepts of *state predicate*, *trajectory*, *number of hops (of a trajectory)*, *non-blocking*, *non-zeno*.

We, however, want to emphasize that to simplify narration, in the following, unless explicitly denoted, “model checking” refers to “*model checking of finite-horizon reachability semantics*”, i.e., whether a state σ of hybrid automaton A satisfies $\varphi_1 \exists \mathcal{U}_{\leq T} \varphi_2$, where φ_1 and φ_2 are state predicates of A , and T is the *finite-horizon*. Also, unless explicitly denoted, *we only discuss non-blocking hybrid automata*.

3 HYBRID SYSTEMS MODELING APPROACH

In this section, we shall use laser tracheotomy, a representative MDPnP application [7][10], as the context to discuss the proper hybrid systems modeling approach for MDPnP. We shall see through this case study why offline model checking must be replaced by online model checking.

Laser tracheotomy MDPnP interlocks various medical devices to increase safety. It has the following entities (see Fig. 1):

- Patient:** the patient that receives the surgery;
- O_2 Sensor:** the patient’s trachea oxygen level sensor;
- SpO_2 Sensor:** the patient’s blood oxygen level sensor;
- Ventilator:** the medical device that administrates the patient’s respirations;
- Surgeon:** the doctor that conducts the surgery;
- Laser Scalpel:** the medical device for the surgeon to cut the patient’s trachea;
- Supervisor:** the central computer that connects all medical devices and makes decisions to guarantee safety.

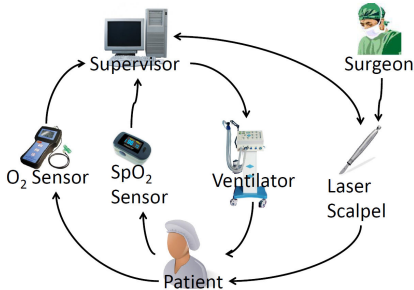


Fig. 1. Layout of Laser Tracheotomy MDPnP

The application context is as follows. In the surgery, due to general anesthesia, the patient is paralyzed, hence has to depend on the ventilator to breathe. The ventilator has three modes: pumping out (the patient inhales oxygen), pumping in (the patient exhales), and hold (the patient exhales naturally due to chest weight). However, when the laser scalpel is to cut the patient’s trachea, the oxygen level inside the trachea must be lower than a threshold. Otherwise, the laser may trigger fire. Therefore, before the laser scalpel is allowed to emit laser, the ventilator must have stopped pumping out (oxygen) for a while. On the other hand, the ventilator can neither stop pumping out for too long, or the patient will suffocate due to too low blood oxygen level.

In summary, the laser tracheotomy MDPnP must avoid the following safety hazards:

Safety Hazard 1: when the laser scalpel emits laser, the patient’s trachea oxygen level exceeds a threshold Θ_{O_2} ;

Safety Hazard 2: the patient’s blood oxygen level reaches below a threshold Θ_{SpO_2} .

Note that the setting of constant thresholds Θ_{O_2} and Θ_{SpO_2} are medical experts’ responsibility and are beyond the coverage of this paper.

The formal expressions of safety hazards will become clear by the end of Section 3.2, when the corresponding hybrid automata are defined.

3.1 Traditional Approach: Offline Modeling

Because the laser tracheotomy MDPnP involves both discrete medical device logic and physical world patient, it is a hybrid system. Therefore we try to model laser tracheotomy MDPnP with hybrid automata.

The traditional approach of model checking, including hybrid systems model checking, is carried out offline. That is, the model is built and its infinite-horizon behavior is verified *before* the system runs. We choose to start with this approach. As a common practice, our offline modeling of laser tracheotomy MDPnP assumes a global time t : t is initialized to 0 second, and $\dot{t} \equiv 1$.

Intuitively, we intend to start with modeling the patient, the core entity of the laser tracheotomy MDPnP. However, the patient’s behavior is directly administrated by the ventilator, which has to be understood first.

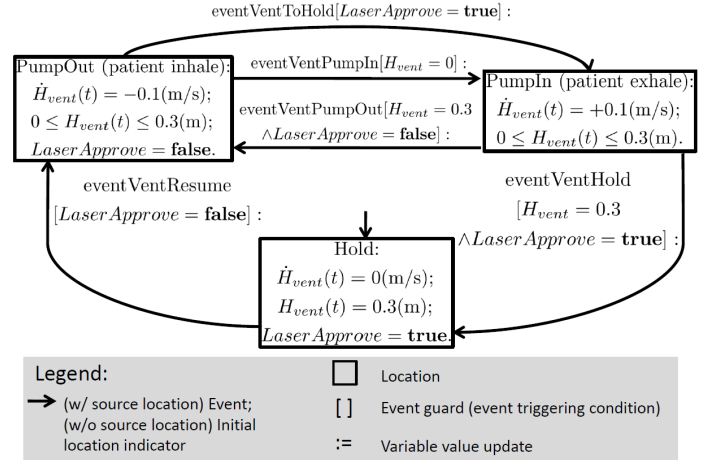


Fig. 2. Offline hybrid automaton of Ventilator

The ventilator is basically a compressible air reservoir [17]: a cylinder of height $H_{vent}(t)$ ($0 \leq H_{vent}(t) \leq 0.3(m)$). The movement of the ventilator cylinder (indicated by $\dot{H}_{vent}(t)$) pumps out/in oxygen/air to/from patient, thus helping the patient to inhale/exhale. The ventilator behavior is defined by the hybrid automaton in Fig. 2. The automaton has three locations: PumpOut, PumpIn, and Hold. When the supervisor (will be discussed later in Fig. 8) allows the ventilator to work (i.e., when data variable *Laser Approve* is set to **false**), the ventilator switches between pumping out (where $\dot{H}_{vent} = -0.1(m/s)$) and pumping in (where $\dot{H}_{vent} = +0.1(m/s)$). This causes the patient to inhale oxygen and exhale respectively. When the

supervisor pauses the ventilator (i.e., when *LaserApprove* is set to **true**), the ventilator cylinder will try to restore to its maximum height (0.3m) and holds there until the ventilator is allowed again (*LaserApprove* set to **false**).

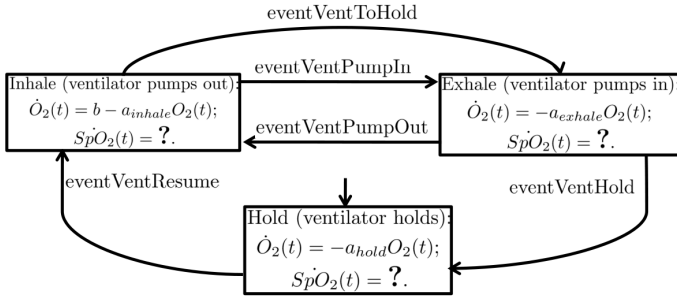


Fig. 3. Offline hybrid automaton of Patient. Though good offline models for \dot{O}_2 exists [7], the offline model for $Sp\dot{O}_2$ is still an open problem. Also note that in location Hold (which corresponds to ventilator Hold), the patient still exhale due to chest weight.

With the ventilator hybrid automaton at hand, we can now start modeling the patient. The patient hybrid automaton (see Fig. 3) is tightly coupled with the ventilator hybrid automaton (see Fig. 2). It also has three locations: Inhale, Exhale, and Hold, which respectively correspond to the ventilator hybrid automaton's locations of PumpOut, PumpIn, and Hold. The events between the three locations are also triggered by corresponding events from the ventilator hybrid automaton.

Inside of each location are the offline continuous time models for trachea oxygen level $O_2(t)$ and blood oxygen level $SpO_2(t)$. Unfortunately, though there are good offline models for $\dot{O}_2(t)$ [7], the offline model for $Sp\dot{O}_2(t)$ is still an open problem [8][9]. This is because blood oxygen level are strongly affected by complex human body biochemical reactions, even emotions.

Therefore, we fail to model $Sp\dot{O}_2(t)$ offline, and hence fail to model the patient offline. What is worse, as the patient model is an indispensable component of the holistic offline model, *the offline model checking of laser tracheotomy MDPnP fails*.

3.2 Proposed Approach: Online Modeling

The failure of offline approach forces us to consider the proposed online approach (see Section 1) instead. Specifically, we sample the patient's trachea/blood oxygen level every T seconds. Suppose at $t_0 = kT$ ($k \in \mathbb{Z}_{\geq 0}$), we get the most up-to-date trachea/blood oxygen level sensor reading $\widehat{O}_2(t_0)$ and $\widehat{SpO}_2(t_0)$, we can then build the hybrid systems model for interval $[t_0, t_0 + T]$, where T is therefore the *finite-horizon*. This model is built as follows.

First, same as the offline model checking, we use global variable t to represent the global clock, except that now t is initialized to t_0 and stops at $(t_0 + T)$ as we only care about the system's finite-horizon safety until $(t_0 + T)$.

The patient hybrid automaton now looks like Fig. 4(a). The biggest change is the continuous time model for the blood oxygen level $Sp\dot{O}_2(t)$. In offline model checking, we have to describe the infinite-horizon behavior of $Sp\dot{O}_2(t)$, which is an

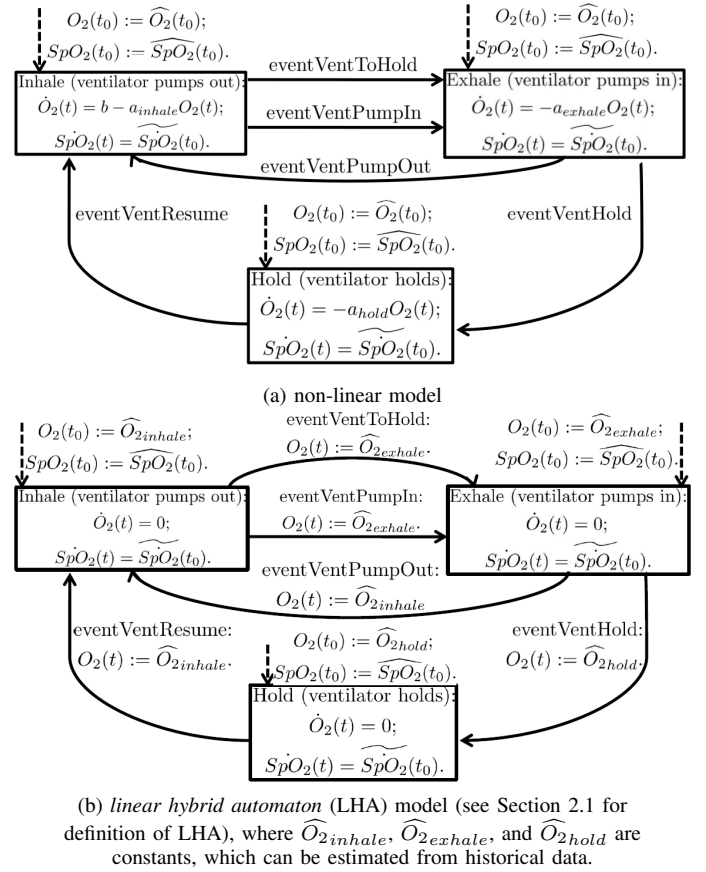


Fig. 4. Online hybrid automaton of Patient.

open problem. However, in online model checking, we only have to describe $Sp\dot{O}_2(t)$'s behavior in interval $[t_0, t_0 + T]$, where the finite-horizon T is just a few seconds. If we only look into such short-run future, blood oxygen level curve $SpO_2(t)$ is very describable and predictable. For example, it cannot plunge from 100% to 10% within just 4 seconds, neither can it show a saw-toothed wave form. Instead, it must be smooth; in fact smooth enough to be safely predicted with standard tools (such as linear regression) based on its past history.

In Fig. 4(a), we use a simple way to predict/describe $Sp\dot{O}_2(t)$ in $t \in [t_0, t_0 + T]$:

$$\dot{SpO}_2(t) \equiv \widetilde{Sp\dot{O}_2}(t_0), \quad \forall t \in [t_0, t_0 + T],$$

where $\dot{SpO}_2(t)$ is the derivative of $SpO_2(t)$ at time t ; and $\widetilde{Sp\dot{O}_2}(t_0)$ is the estimation (e.g., via linear regression) of $\dot{SpO}_2(t_0)$ based on $SpO_2(t)$'s history recorded during $(t_0 - T_{past}, t_0)$. T_{past} is a configuration constant picked empirically offline. In our case study, we pick $T_{past} = 6$ seconds.

Also, depending on the patient's state at time t_0 , the initial location can be Inhale, Exhale, or Hold. Whichever location it is, the initial value of trachea/blood oxygen value should be $\widehat{O}_2(t_0)$ and $\widehat{SpO}_2(t_0)$ respectively.

The patient model of Fig. 4(a) can be further simplified. Human subject respiration traces (see Fig. 5) show that the values of a_{inhale} , a_{exhale} , and a_{hold} in Fig. 4(a) are large: so large that $O_2(t)$ almost behaves as rectangular waves when

the patient hybrid automaton changes locations. Therefore, we can simplify Fig. 4(a) into Fig. 4(b), where $O_2(t)$ remains constant within every location, and its value is only updated on the corresponding transitions. This simplification turns the patient hybrid automaton (in fact the whole system) into an *linear hybrid automaton* (LHA) (see Section 2.1 for definition of LHA), which is much easier to verify [18].




Fig. 5. A typical example excerpt of trachea CO_2 level trace (measured on human subjects with Nonin 9843 [19]); note $O_2(t) = C_1 - C_2 \cdot CO_2(t)$, where C_1 and C_2 are two constants, whose derivation can be found in classic physics textbooks [20].

We now check other laser tracheotomy MDPnP entities.

First, since the online model only looks into the finite-horizon of $[t_0, t_0 + T]$, where T is also the sensor sampling period, there are no interactions with sensors throughout the interval of $(t_0, t_0 + T)$. Therefore, in online model checking, the hybrid automata of O_2 sensor and SpO_2 sensor are unnecessary.

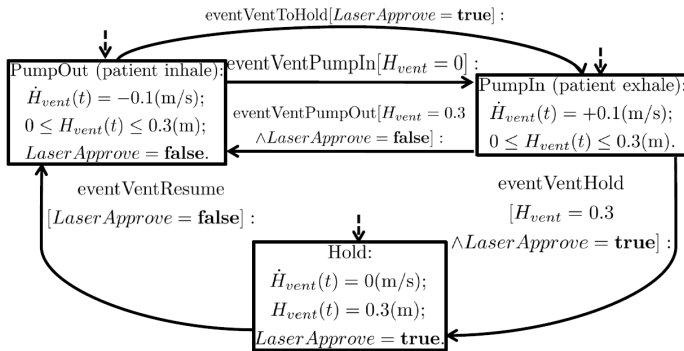


Fig. 6. Online hybrid automaton of Ventilator.

Next, the ventilator hybrid automaton in online model (see Fig. 6) is almost the same as its offline counterpart (see Fig. 2) A main difference is that the online model's initial location can be any location depending on the ventilator's state at t_0 .

The last entity that directly interacts with the patient is the laser scalpel. We can actually model the laser scalpel and the surgeon with one hybrid automaton: the laser scalpel hybrid automaton (see Fig. 7).

The automaton's key elements are the two Boolean variables: $LaserApprove$ and $LaserReq$.

$LaserApprove$ indicates whether the supervisor (see Fig. 1) allows the laser scalpel to emit laser (**true** for yes and **false** for no). Its value can only be set by the supervisor hybrid automaton (see Fig. 8), which is to be explained later.

$LaserReq$ indicates whether the laser scalpel wants to emit laser (**true** for yes and **false** for no). Its value can only be set by the laser scalpel hybrid automaton. The value setting is

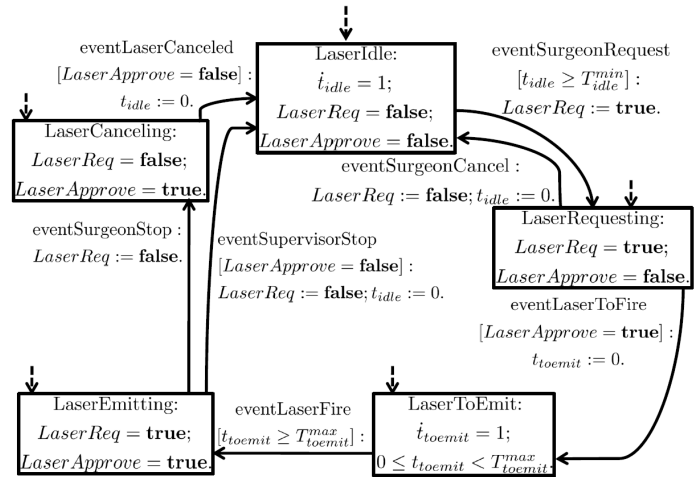


Fig. 7. Online hybrid automaton of Laser Scalpel. This is the only automaton that sets the value of state variable $LaserReq$.

triggered by following events: *i*) when in $LaserIdle$, the surgeon can request emitting laser through $eventSurgeonRequest$, which sets $LaserReq$ to **true**; *ii*) when in $LaserRequesting$ or $LaserEmitting$, the surgeon can request stopping laser emission through $eventSurgeonCancel$ and $eventSurgeonStop$ respectively, which both set $LaserReq$ to **false**; *iii*) when in $LaserEmitting$, the supervisor can stop the laser emission at any time by setting $LaserApprove$ to **false**, which triggers $eventSupervisorStop$ and sets $LaserReq$ to **false**.

The four possible combinations of $LaserApprove$ and $LaserReq$'s values define the major locations in the laser scalpel hybrid automaton: $LaserIdle$, $LaserRequesting$, $LaserEmitting$, and $LaserToEmit$. Particularly, laser scalpel emits laser in and only in $LaserEmitting$. There is an additional location, $LaserToEmit$, which models the additional delay T_{toemit}^{max} between $LaserRequesting$ and $LaserEmitting$. This delay is to further ensure oxygen level in trachea falls below threshold before the actual laser emission.

The laser scalpel hybrid automaton's initial location can be anywhere depending on the laser scalpel's state at t_0 . One thing to note is that all variables should be initialized to their actual value at t_0 . For example, if initial location is $LaserIdle$, and Laser Scalpel has been idling for 10 seconds by t_0 , then t_{idle} shall be initialized to 10 seconds instead of 0.

Finally, all medical device entities are interlocked by the supervisor, the central decision making computer (see Fig. 1). The supervisor maneuvers data variable $LaserApprove$. Setting $LaserApprove$ to **true/false** determines the off/on of the ventilator and the permission/denial of emitting laser respectively.

The value setting decisions are made dependent on the most up-to-date information on the patient's trachea oxygen level $O_2(t)$ and blood oxygen level $SpO_2(t)$. Based on the models given in the patient hybrid automaton (see Fig. 4), we can predict $O_2(t)$ and $SpO_2(t)$ for any $t \in [t_0, t_0 + T]$. Therefore, we can construct the supervisor hybrid automaton as Fig. 8, which directly uses $O_2(t)$ and $SpO_2(t)$ predicted by the patient hybrid automaton for decision making.

The supervisor hybrid automaton has two locations:

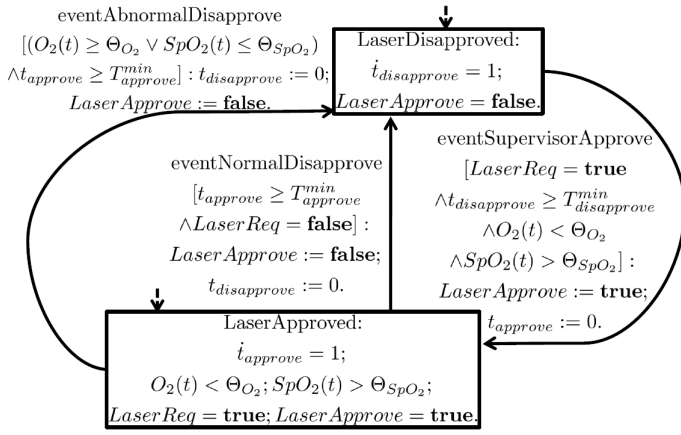


Fig. 8. Online hybrid automaton of Supervisor. This is the only automaton that sets the value of data variable $LaserApprove$. Note $t_{approve}$ can be totally removed from the model in *soft* real-time online model checking.

LaserDisapproved and LaserApproved.

When in LaserDisapproved, the supervisor needs eventSupervisorApprove to move to LaserApproved. This event is triggered when the following prerequisites all hold:

Prerequisite 1: the laser scalpel is requesting emitting laser (i.e., $LaserReq = true$);

Prerequisite 2: $O_2(t)$ is less than threshold Θ_{O_2} ;

Prerequisite 3: $SpO_2(t)$ is greater than threshold Θ_{SpO_2} .

Prerequisite 4: $t_{disapprove} \geq T_{disapprove}^{min}$. This is a minimal dwelling time requirement to guarantee the automaton’s non-zero property. The purpose will become clear in a latter example (Example 1 of Appendix B in the supplementary file). This requirement also models the time cost in switching between LaserDisapproved and LaserApproved modes in the supervisor.

Through eventSupervisorApprove, the supervisor approves the emission of laser by setting $LaserApprove$ to **true**. This event also resets a clock $t_{approve}$, and moves the location to LaserApproved.

Like $t_{disapprove}$, clock $t_{approve}$ is for guaranteeing a minimal dwelling time of $T_{approve}^{min}$ in LaserApproved. After that, if Prerequisite 1 no longer holds (i.e., when $LaserReq$ becomes **false**), the eventNormalDisapprove is triggered. This event moves the supervisor back to location LaserDisapproved and resets $LaserApprove$ to **false**, and $t_{disapprove}$ to 0.

In contrast to eventNormalDisapprove, eventAbnormalDisapprove is triggered when the supervisor is in LaserApproved while Prerequisite 2 or 3 stops to hold. This event also moves the supervisor back to location LaserDisapproved and resets $LaserApprove/t_{disapprove}$ to **false/0** respectively.

Finally, same as the other online hybrid automata, the initial location for the online supervisor automaton can be either LaserDisapproved or LaserApproved, depending on the state of the supervisor at time t_0 ; and the variables should be initialized to the actual values at t_0 .

With the above hybrid automata model of the laser tracheotomy MDPnP, we can formally express Safety Hazard 1 and 2 (see the beginning of Section 3) as follows.

Safety Hazard 1: For any given initial state σ_0 , $\sigma_0 \models true \exists \mathcal{U}_{\leq T} \bigcup_{v \in V_{comp} \wedge v|_{ls} = LaserEmitting} (v, O_2(t) \geq \Theta_{O_2})$;

Safety Hazard 2: For any given initial state σ_0 , $\sigma_0 \models true \exists \mathcal{U}_{\leq T} \bigcup_{v \in V_{comp}} (v, SpO_2(t) \leq \Theta_{SpO_2})$;

where V_{comp} is the location set of the combined automaton of the Ventilator, Patient, Laser Scalpel, and Supervisor; $v|_{ls}$ is v ’s projection on the Laser Scalpel automaton location set.

When model checking any one of the above safety hazards, a “yes” answer means the system is unsafe; while a “no” answer means this system is safe.

4 SYSTEM CO-DESIGN PATTERN

The evolution from offline model checking to online model checking must also be matched with system design changes.

4.1 Hard Real-Time System Design

First, the overall system architecture shall integrate online model checking as a runtime fault prediction and prevention mechanism.

A straightforward thought is to run online model checking periodically. So far, we have assumed the period to be the same as the online model checking’s finite-horizon T . That is, at the beginning of each period T , online model checking predicts whether unsafe states are reachable within the coming T seconds. If so, the system switches to a fall-back plan for the current period. The fall-back plan is application dependent. For laser tracheotomy MDPnP, a simple fall-back plan is that the supervisor locks $LaserApprove$ at **false**, hence forbidding laser emission and keeping the ventilator active.

The above overall architecture works if online model checking costs 0 time. In practice, this is an over simplification. However, if the online model checking has a *worst case execution time bound* $D < T$ (where T is the online model checking’s finite-horizon), then we can run the online model checking as a *hard real-time* task and use pipelining to carry out fault prediction and prevention. This is formally described by the algorithm in Fig. 9, which, without loss of generality, runs a pipeline with $T = 2D$; and D replaces T to be the new sampling period.

```
//This code assumes online model checking (see line 4, 5) can always
//finish within hard real-time deadline  $D = \frac{T}{2}$ .
1. main(){
2.   wait till current time  $t$  satisfies  $(t \bmod \frac{T}{2} = 0)$ ;
3.    $t_0 := t$ ;
4.   read sensors and build online model  $A$ ;
5.   if ( $A$  may reach unsafe states in  $[t_0, t_0 + T]$ ){
6.     /*non-blocking call:*/ switch the hybrid system to fall-back plan;
7.   }else
8.     /*non-blocking call:*/ allow the hybrid system to run normally;
9.   goto line 2;
```

Fig. 9. Overall system architecture for *hard real-time* online model checking, with worst case execution time bound of D (for line 4, 5). Without loss of generality, the code runs a pipeline with $T = 2D$ (see line 2, 5). To “run normally” means that the hybrid system runs according to online model A ’s (see line 4) descriptions.

To run the hard real-time algorithm of Fig. 9, the online model checking problem must be *decidable*. That is, a time

cost upper bound must exist. In the following, we show a large family of hybrid automata systems, *strongly non-zeno LHA systems* (SNZ-LHA-Systems) [21] to be exact, satisfy the decidability requirement.

Definition 1 (SNZ-LHA-System): Let \mathcal{S} be a set of *linear hybrid automata* (LHA). For each LHA $A \in \mathcal{S}$, let $\mathcal{T}_A \stackrel{\text{def}}{=} \{\tau | \tau \text{ is a trajectory (see [15] for the definition of "trajectory") of } A \text{ and } \tau \text{ passes a transition of } A \text{ twice}\}$. If $\exists \varepsilon > 0$, such that $\forall A \in \mathcal{S}, \inf_{\tau \in \mathcal{T}_A} \{\delta_\tau\} \geq \varepsilon$ (where δ_τ is τ 's duration; $\inf \emptyset \stackrel{\text{def}}{=} \infty$), then \mathcal{S} is called a *strongly non-zeno LHA system* (SNZ-LHA-System).

For an SNZ-LHA-System, we have the following:

Theorem 1 (Decidability): Finite-horizon reachability model checking of an SNZ-LHA-System is decidable.

Proof: See Appendix A in the supplementary file. ■

What is more, the proof of Theorem 1 also shows a time cost upper bound for finite-horizon reachability model checking of an SNZ-LHA-system exists. In fact, interested readers can refer to [22] for a loose time cost upper bound, though a tight time cost upper bound is still an open problem.

Therefore, if we ensure an online hybrid systems model to be an SNZ-LHA-System, *real-time* worst case execution time (i.e., deadline) exists.

Given a set \mathcal{S} of LHAs, we claim in the following that \mathcal{S} is ensured to be an SNZ-LHA-System if it complies with certain design patterns stated in Theorem 2.

Theorem 2 (Decidable Design Pattern): If every cycle of transitions in \mathcal{S} complies with one of the following design patterns: ε -Minimal Dwelling Time, ε -Alternating Cyber-Value, or ε -Alternating Physical-Value, then finite-horizon reachability model checking on the LHA set \mathcal{S} is decidable.

Proof: See Appendix B in the supplementary file for detailed definitions and proof. ■

If we review the laser tracheotomy MDPnP online LHA model (see Fig. 4(b) ~ 8), we find its design pattern complies with Theorem 2. Hence *online finite-horizon reachability hybrid systems model checking* (simplified as "*online model checking*" in the following, unless explicitly denoted) on laser tracheotomy MDPnP is decidable. That is, theoretically, a worst case execution time bound for hard real-time exists.

4.2 Soft Real-Time System Design

Though online *hard* real-time model checking of SNZ-LHA-Systems is theoretically possible due to Theorem 1, a tight bound on worst case execution time is still an open problem. A very loose bound is known (see [22]), but it is often too large to be practical. In fact, we know the following:

Theorem 3: Finite-horizon reachability model checking of an SNZ-LHA-System is NP-Hard.

Proof: See Appendix C in the supplementary file. ■

Theorem 3 implies online hard real-time model checking of SNZ-LHA-Systems is only practical for very small scale

cases; *soft* real-time online model checking instead has more practical value.

In soft real-time online model checking, we directly specify a desired deadline D , without requiring *hard* real-time guarantee. The selection method of D is empirical: as long as D makes deadline misses satisfactorily rare and the online modeling satisfactorily accurate. For example, we can use standard benchmarks to find a desirable D (see Section 5.2).

Even though deadline D may be missed, soft real-time online model checking can still serve the MDPnP hybrid system in at least two ways: one conservative and the other aggressive, as described by the pseudo code in Fig. 10.

```
//Online model checking deadline is  $D = \frac{T}{2}$  (see line 4, 6, 7, 11, 12).
1. main(mode){
2.   wait till current time  $t$  satisfies  $(t \bmod \frac{T}{2} = 0)$ ;
3.    $t_0 := t$ ;
4.   read sensors and build online model  $A$ ;
5.   if (mode = "conservative way"){
6.     if (( $A$  may reach unsafe states in  $[t_0, t_0 + T]$ )
7.       or (current time  $t \geq t_0 + \frac{T}{2}$ )){
8.       /*non-blocking call:*/ switch the hybrid system to fall-back plan;
9.     }else
10.    /*non-blocking call:*/ allow the hybrid system to run normally;
11.   }else {//mode = "aggressive way"
12.     if ((not ( $A$  may reach unsafe states in  $[t_0, t_0 + T]$ )
13.       or (current time  $t \geq t_0 + \frac{T}{2}$ )){
14.       /*non-blocking call:*/ allow the hybrid system to run normally;
15.     }else
16.     /*non-blocking call:*/ switch the hybrid system to fall-back plan;
17.   }
```

Fig. 10. Revised overall system architecture that allows *soft real-time* online model checking. Without loss of generality, the code runs a pipeline with $T = 2D$ (see line 2, 6, 11), where $D = \frac{T}{2}$ is the real-time online model checking deadline. To "run normally" means that the hybrid system runs according to online model A 's (see line 4) descriptions.

In the conservative way, if online model checking misses deadline D , the MDPnP hybrid system always switches to the (application dependent) fall-back plan. Assuming the modeling is accurate, the conservative way can prevent all accidents. However, if deadline misses are too often, the system will frequently switch to fall-back plan, annoying the users. In other words, the conservative way can raise a lots of false alarms, but can prevent all accidents.

Take our laser tracheotomy MDPnP for example. Every time the online model checking misses the D seconds deadline on safety check, the supervisor will disapprove any laser emission request for the next D seconds (i.e., the "fall-back plan"). Instead, only when the online model checking confirms safety within the D seconds deadline will the supervisor follow Fig. 8's descriptions in the next D seconds.

In the aggressive way, if online model checking misses deadline D , the MDPnP system does not switch to fall-back plan. The aggressive way only invokes fall-back plan when it is certain the system is facing risks. In other words, the aim of aggressive way is not to prevent all accidents, but to *reduce* accidents. In medical practice, a method that can significantly reduce accidents is still a useful method; in fact, most medical routines are of such nature [23].

Again take our laser tracheotomy MDPnP for example. Every time the online model checking misses the D seconds deadline on safety check, the supervisor will nevertheless follow Fig. 8’s descriptions in the next D seconds. The fall-back plan (that the supervisor disapproves any laser emission requests) only kicks in when online model checking is certain that unsafe state is reachable within the D seconds deadline. Therefore, the online model checking is not to *eliminate* all possible accidents that a human surgeon may make, but to *reduce* such accidents as an additional protection.

To summarize, each deadline miss means the online model checking is uncertain about the safety of the MDPnP hybrid system in the next D seconds. In the conservative way, the system always switches to the fall-back plan when the online model checking ends up uncertain (of course it also switches to the fall-back plan when the online model checking is certain of pending risks). In the aggressive way, the system only switches to fall-back plan when the online model checking is certain of pending risks.

5 EVALUATIONS

To validate our proposed approach, especially the effectiveness (usefulness) of soft real-time online model checking for MDPnP (the “conservative way” and the “aggressive way”, see Section 4.2), we carry out evaluations using real-world trachea/blood oxygen level traces.

5.1 Effectiveness

We run soft real-time online model checking program \mathcal{P} (see Fig. 10) upon emulated trachea/blood oxygen level sensors for 1200 seconds. We choose soft real-time deadline to be $D = 2$ seconds (see Section 5.2 for why). According to the soft real-time pseudo code of Fig. 10, this means every $D = \frac{T}{2} = 2$ seconds, \mathcal{P} queries the emulated sensors for trachea/blood oxygen level readings, then builds and verifies an online model with finite-horizon of $T = 2D = 4$ seconds.

We have two sets of 1200-second traces for the emulated sensors. The first set of 1200-second traces comes from PhysioNet [24], a comprehensive online public database (set up by NIH, NIBIB, and NIGMS) of real-world medical traces logged by hospitals. For simplicity, we call it “PhysioNet Traces”. The other set of 1200-second traces comes from our own experiments on two human subjects. Human Subject 1 (HS1) mimics the combined behavior of the supervisor, laser scalpel, and surgeon in laser tracheotomy MDPnP. As shown by Fig. 11(a), HS1 randomly swaps between holding the flag of “Laser Disapproved” and “Laser Approved”. Human Subject 2 (HS2) mimics the combined behavior of the ventilator and the patient in the laser tracheotomy MDPnP. When HS1 holds the “Laser Disapproved” flag, HS2 breathes smoothly at the rate of 6 seconds per respiration-cycle. When HS1 holds the “Laser Approved” flag, HS2 first tries to exhale (to his very best) and then holds his breath until HS1 raises the “Laser Disapproved” flag again (in case HS1 holds the “Laser Approved” flag for too long, HS2 is free to abort the experiment by resuming normal breath). Meanwhile, HS2’s trachea and blood oxygen




Fig. 11. Human subjects roles and behaviors. (a) HS1; (b) HS2.

level are recorded by Nonin 9843 [19]. We call the derived traces the “HKPolyU Traces”.

The two emulated sensors read corresponding real-world traces (PhysioNet or HKPolyU) respectively. Based on the readings, \mathcal{P} builds online hybrid systems models as described in Section 3.2, and verifies it. The specific modeling and verification software used is PHAVer [18], a well-known hybrid systems model checking tool. Our computation platform is a Lenovo Thinkpad X201 with Intel Core i5 and 2.9G memory; the OS is 32-bit Ubuntu 10.10.

For each trace, throughout its 1200-second emulation period, program \mathcal{P} carries out $1200/D = 1200/2 = 600$ trials of online modeling and verifications. The statistics of execution time cost is depicted by Table 1.

The statistics show that more than 97.8% of the online model checking trials finished within the $D = 2$ (sec) deadline. In other words, only no more than 2.2% of the online model checking trials missed deadline.

Assume the modeling is accurate (which is going to be validated soon), in case \mathcal{P} runs the “conservative way” (see Fig. 10), the above result means not only all accidents are prevented, the false alarm probability is no more than 2.2%. In case \mathcal{P} runs the “aggressive way”, the above result means more than 97.8% of accidents can be reduced (every time the system can reach unsafe states in the next D seconds, there is a $\geq 97.8\%$ chance that online model checking finishes within deadline, hence triggering the fall-back plan). Such reduction of accidents is significant according to the standards of medical practice [23]. In either case, the results provide strong evidence that (soft) real-time online model checking is effective (i.e., feasible and useful).

TABLE 1

Statistics of execution time cost of online model checking (unit: second; deadline $D = 2$ seconds)

| | % of trials missed deadline | Execution time of those caught deadline (secs) | | | |
|-----------------|-----------------------------|--|-------|-------|-------|
| | | Min | Max | Mean | Std |
| PhysioNet Trace | 2.2% | 0.817 | 1.720 | 0.932 | 0.126 |
| HKPolyU Trace | 1.7% | 0.818 | 1.940 | 0.965 | 0.146 |

To validate the assumption that the online modeling is accurate, we carry out statistics on the prediction error of blood

oxygen level curve.

During the online model checking, at every time instance $t_0 = kD$ ($k \in \{0, 1, \dots, 599\}$, and $D = 2$ seconds), we sample the blood oxygen level and predict (see Fig. 4) the blood oxygen level curve in $[t_0, t_0 + T]$ ($T = 2D = 4$ seconds). Let the predicted blood oxygen level at time $(t_0 + T)$ be $\widetilde{SpO}_2(t_0 + T)$. Let the PhysioNet/HKPolyU trace reading of blood oxygen level at time $(t_0 + T)$ be $\widehat{SpO}_2(t_0 + T)$. We define the relative prediction error at time $(t_0 + T)$ to be

$$ERR_{SpO_2}(t_0 + T) = \frac{|\widetilde{SpO}_2(t_0 + T) - \widehat{SpO}_2(t_0 + T)|}{\widehat{SpO}_2(t_0 + T)}.$$

The statistics of the relative prediction errors throughout the 600 trials for each trace are depicted by Table 2. The statistics show that our online model checking’s predictions on the finite-horizon behavior of blood oxygen level curve match the real-world traces quite accurately (with maximum relative error of 3.92%).

TABLE 2

Statistics of blood oxygen level online modeling relative errors (%)

| | Min | Max | Mean | Std |
|-----------------|--------|------|------|------|
| PhysioNet Trace | 0.03 | 2.53 | 0.51 | 0.52 |
| HKPolyU Trace | < 0.01 | 3.92 | 0.61 | 0.60 |

5.2 Selection of Soft Real-Time Online Model Checking Deadline

Now we show why $D = 2$ seconds is an empirically desirable soft real-time online model checking deadline for the pseudo code of Fig. 10.

We use both the 1200-second PhysioNet Trace and the 1200-second HKPolyU Trace as benchmark, and try out different values of D .

Table 3 shows the statistics on online modeling relative errors under different D s. The statistics show that $D = 2$ seconds incurs least maximum relative error compared to other candidates. Note $D = 2$ seconds might not be the optimal choice, but based on the evaluations on the 2400-second medical traces, it turns out to be an empirically effective choice. A lot of parameters used in medicine are derived from such empirical studies.

TABLE 3

Online Model Checking Relative Error Statistics under Different D s

| Trace | D (sec) | Relative Error (%) | | | |
|-----------|-----------|--------------------|------|------|------|
| | | Min | Max | Mean | Std |
| PhysioNet | 2 | 0.03 | 2.53 | 0.51 | 0.52 |
| | 3 | 0.04 | 4.52 | 0.76 | 0.74 |
| | 4 | < 0.01 | 5.98 | 0.96 | 0.94 |
| HKPolyU | 2 | < 0.01 | 3.92 | 0.61 | 0.60 |
| | 3 | < 0.01 | 4.81 | 0.90 | 0.90 |
| | 4 | < 0.01 | 6.29 | 1.18 | 1.12 |

6 DISCUSSIONS

6.1 False Negatives and False Positives

If the online model is absolutely accurate, the online model checking either misses deadline, or produces true-positive/true-negative conclusions.

Interestingly, even if the online model is inaccurate, i.e., if the online model checking *can* produce false-positive/false-negative conclusions, our proposed method can still be useful for medical practices. Please see Appendix D in the supplementary file for details.

6.2 Wireless Communications Links

So far, we have assumed reliable communications links between entities. Though this assumption is empirically valid for wired communications links, it is not for wireless.

How to adopt unreliable wireless communications links in life/safety critical medical settings is a nontrivial and active research area [25][26][27][28]. A comprehensive solution is beyond the scope of this paper. However, we can still provide a simple hybrid solution to allow wireless links between the sensors and the supervisor. Our solution is as follows.

According to the pseudo code of Fig. 10, every D seconds, the sensors are supposed to update the supervisor with the new readings of the patient’s vital sign(s). Suppose at time instance iD ($i \in \mathbb{Z}_{\geq 0}$), the corresponding reading is X_i . Suppose at time instance iD , the supervisor needs to look at X_{i-k} , X_{i-k+1} , \dots , X_i to build the online model. If any reading(s) of $X_{i-k} \sim X_i$ is(are) lost due to wireless communications failures, then for the period of $[iD, (i+1)D]$, the supervisor shall refuse to carry out online model checking, to cause a deliberate “deadline miss”. This deliberately created deadline miss shall then be treated as a usual deadline miss.

In this way, any wireless communications failures will only result in more deadline misses. The designs and analysis described in the previous sections (and subsections) still sustain.

For further evaluations of this wireless approach, please refer to Appendix E of the supplementary file.

7 RELATED WORK

Our approach is different from the well-known runtime verification [29]. Runtime verification aims to discover latent bugs of programs by logging and analyzing the programs’ execution traces under varied inputs/configurations. It is not for *predicting/preventing* faults *before* they ever happen; whilst our approach is. For many medical CPS systems, the cost/consequence of possible faults in test runs is high or even unbearable. This necessitates our approach of predicting and *preventing* faults before they ever happen.

Sen et al. [30] propose an online safety analysis method for multithreaded programs. However, this work only focuses on how to infer other potential executions that can take place in the *past*. Our work tries to predict the *future* state of patient based on recent observations

Easwaran et al. [31], Qi et al. [32], and Harel et al. [33] also propose bringing model checking online. But they are still focusing on discrete (automata) model checking, rather than hybrid systems model checking that this paper is about.

Sauter et al. [34] propose a lightweight hybrid-system model checking method, which uses ordinary differential equations (ODE) to predict temporal logic properties. However, in the MDPnP systems it is not uncommon to be lack of differential equations governing patients dynamics, i.e., patients model.

Li et al. [35] propose one online model checking approach aiming at automatically estimating parameters in simulation models, which are often used for biological purpose to understand complex regulatory mechanisms in cell.

Larsen et al. [36] propose an online model-based testing tool for real-time systems, UPPAAL TRON. The tool is based on UPPAAL engine and models real-time systems as timed automata, whereas our online model checking of MDPnP systems focuses on more general hybrid systems.

Also, our approach is not model-checker specific, though our evaluation in this paper uses PHAVer. In fact, we are considering integrating our approach with other well-known model checkers, such as Bogor [37], CellExcite [38] etc..

8 CONCLUSIONS AND FUTURE WORK

Through our case study on laser tracheotomy MDPnP, we show that online model checking of short-run future behavior can effectively address the two challenges in MDPnP CPS hybrid systems model checking. By focusing on online and short-run future, many originally hard to describe/predict human body parameters become describable and predictable; and many variable parameters become fixed numerical values, which greatly simplifies verification. The online model checking can go real-time if the proposed hard/soft real-time system co-design patterns are followed. Our empirical evaluations based on real-world human subject traces show that our online model checking and co-design approach is feasible and effective. As future work, we will carry out more evaluations and integrate/extend our approach to more comprehensive MDPnP/CPS frameworks [39][40][41][42].

REFERENCES

- [1] *Medical Device Plug-and-Play (MDPnP)*. <http://www.mdnp.org>.
- [2] L. Sha et al., "Cyber-physical systems: A new frontier," *Machine Learning in Cyber Trust: Security, Privacy, and Reliability*, 2009.
- [3] C. Baier et al., *Principles of Model Checking*. MIT Press, 2008.
- [4] P. Tabuada, *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, 2009.
- [5] P. J. Antsaklis and X. D. Koutsoukos, "Hybrid systems: Review and recent progress," *Software-Enabled Control: Information Technology for Dynamical Systems*, pp. 273–298, 2003.
- [6] I. Lee and O. Sokolsky, "Medical cyber physical systems," *DAC*, 2010.
- [7] C. Kim et al., "A framework for the safe interoperability of medical devices in the presence of network failures," *ICCPs'10*, Apr. 2010.
- [8] J. A. Grass, "Patient-controlled analgesia," *Anesthesia & Analgesia*, vol. 101, no. 5S, pp. S44–S61, Nov. 2005.
- [9] J. X. Mazoit, K. Butscher, and K. Samii, "Morphine in postoperative patients: Pharmacokinetics and pharmacodynamics of metabolites," *Anesthesia and Analgesia*, vol. 105, no. 1, pp. 70–78, 2007.
- [10] "Medical devices and medical systems - essential safety requirements for equipment comprising the patient-centric integrated clinical environment (ice), part 1: General requirements and conceptual model," no. STAM F2761-2009, 2009.
- [11] T. Li et al., "From offline toward real-time: A hybrid systems model checking and cps co-design approach for medical device plug-and-play (mdnp)," *Proc. of IEEE/ACM 3rd Intl. Conf. on Cyber-Physical Systems (ICCPs)*, pp. 13–22, Apr. 2012.
- [12] —, "From offline long-run to online short-run: Exploring a new approach of hybrid systems model checking for mdnp," *Joint Workshop on HCMDS/MDPnP*, Apr. 2011.
- [13] L. Bu et al., "Toward online hybrid systems model checking of cyber-physical systems time-bounded short-run behavior," *ICCPs'11 Work-in-Progress Session*, Apr. 2011.
- [14] R. Alur et al., "Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems," *Hybrid Systems*, vol. 736, pp. 209–229, 1992.
- [15] —, "Automatic symbolic verification of embedded systems," *IEEE Trans. on Software Engineering*, vol. 22, no. 3, pp. 181–201, Mar. 1996.
- [16] T. A. Henzinger, P.-H. Ho, and H. Wong-Toi, "Hytech: a model checker for hybrid systems," *STTT*, vol. 1, no. 1-2, pp. 110–122, 1997.
- [17] J. A. Dorsch and S. E. Dorsch, *Understanding Anesthesia Equipment, 5th ed.* Lippincott Williams and Wilkins, 2007.
- [18] G. Frehse, "PHAVer: Algorithmic Verification of Hybrid Systems past HyTech," *Proc. of HSCC'05*, vol. LNCS 2289, pp. 258–273, 2005.
- [19] Nonin 9843 oximeter/Co2 detector. <http://www.nonin.com>.
- [20] M. J. Moran et al., *Fundamentals of Engineering Thermodynamics*. Wiley, 2003.
- [21] O. Roux and V. Rusu, "Uniformity for the decidability of hybrid automata," *Static Analysis*, vol. 1145, pp. 301–316, 1996.
- [22] T. Li et al., *From Offline toward Real-Time: A Hybrid Systems Model Checking and CPS Co-Design Approach for Medical Device Plug-and-Play (Technical Report Appendices)*. <http://www.comp.polyu.edu.hk/~7Ecsqwang/research/appendix.html>.
- [23] F. Brunicaudi et al., *Principles of Surgery*. McGraw-Hill, Sep. 2009.
- [24] *PhysioNet: the Research Resource for Complex Physiologic Signals*. <http://www.physionet.org>.
- [25] Y. Wang et al., "WiCop: Engineering WiFi temporal white-spaces for safe operations of wireless body area networks in medical applications," *Proc. of RTSS'11*, pp. 170–179, 2011.
- [26] J. Huang et al., "Beyond co-existence: Exploiting WiFi white space for ZigBee performance assurance," *Proc. of ICNP*, pp. 305–314, Oct. 2010.
- [27] Q. Wang et al., "Building robust wireless LAN for industrial control with the DSSS-CDMA cell phone network paradigm," *IEEE Trans. on Mobile Computing*, vol. 6, no. 6, pp. 706–719, Jun. 2007.
- [28] S. Baker et al., "Medical-grade, mission-critical wireless networks," *IEEE EMB Magazine*, vol. 27, no. 2, pp. 86–95, 2008.
- [29] B. Finkbeiner et al., "Collecting statistics over runtime executions," *ENTCS*, vol. 70:4, 2002.
- [30] K. Sen, G. Rosu, and G. Agha, "Online efficient predictive safety analysis of multithreaded programs," in *Proceedings of 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'04)*, Barcelona, Spain, 2004, pp. 123–138.
- [31] A. Easwaran et al., "Steering of discrete event systems: Control theory approach," *Workshop on Runtime Verif.*, 2006.
- [32] Z. Qi et al., "A hybrid model checking and runtime monitoring method for c++ web services," *Intl' Joint Conf. on INC, IMS and IDC*, 2009.
- [33] D. Harel, H. Kugler, R. Marelly, and A. Pnueli, "Smart play-out of behavioral requirements," in *Proceedings of the 4th International Conference on Formal Methods in Computer-Aided Design (FMCAD '02)*, London, UK, 2002, pp. 378–398.
- [34] G. Sauterand et al., "Lightweight hybrid model checking facilitating online prediction of temporal properties," in *Proceedings of the 21st Nordic Workshop on Programming Theory, NWPT 09*, Kgs. Lyngby, Denmark, 2009, pp. 20–22.
- [35] C. Li, M. Nagasaki, C. H. Koh, and S. Miyano, "Online model checking approach based parameter estimation to a neuronal fate decision simulation model in caenorhabditis elegans with hybrid functional petri net with extension," *Mol. BioSyst*, pp. 1576–1592, 2011.
- [36] K. G. Larsen, M. Mikucionis, and B. Nielsen, "Uppaal tron user manual," 2007.
- [37] Robby et al., "Bogor: An extensible and highly-modular software model checking framework," *Proc. of the 9th European Software Engineering Conference (ESEC/FSE-11)*, 2003.
- [38] E. Bartocci et al., "Cellexcite: An efficient simulation environment for excitable cells," *BMC Bioinformatics*, vol. 9, no. 2, pp. 1–13, Mar. 2008.
- [39] J. Ko et al., "Wireless sensor networks for healthcare," *Proc. of IEEE*, vol. 98, no. 1, pp. 1947–1960, Nov. 2010.
- [40] A. King et al., "An open test bed for medical device integration and coordination," *Proc. of ICSE'09 (Software Engineering in Practice Track)*, May 2009.
- [41] R. Alur and G. Weiss, "Rtcomposer: a framework for real-time components with scheduling interfaces," *EMSOFT'08*, pp. 159–168, 2008.
- [42] Y. Krishnamurthy, C. Gill et al., "The design and implementation of real-time corba 2.0: Dynamic scheduling in tao," *RTAS*, 2004.



Tao Li Tao Li received the BE degree in software engineering from Southeast University, Nanjing, China, in 2007, the MSc degree in computer science and technology from Nanjing University, Nanjing, China, in 2010. He is currently pursuing the PhD degree in the Department of Computing at the Hong Kong Polytechnic University, Hong Kong. His research interests include cyber-physical systems (CPS), pervasive/ubiquitous computing, and their applications in medical treatment and healthcare. He is

a student member of the IEEE.



Feng Tan Feng Tan received his B.E and M.E degree in Industrial Engineering from Univ. of Electronic Science and Technology of China, Chengdu, China in 2009 and 2012 respectively. He is currently pursuing the Ph.D degree in the department of Computing and the Hong Kong Polytechnic University, Hong Kong. His research interests include Cyber-Physical Systems (CPS), particularly on dependable system architecture design in CPS applications.



Qixin Wang (M'08) received the B.E. and M.E. degrees from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 1999 and 2001, respectively, and the Ph.D. degree from the Department of Computer Science, University of Illinois at Urbana-Champaign (UIUC) in 2008. He was a Research Assistant/Associate in UIUC from 2001 to 2008. Since 2009, he is an Assistant Professor in the Department of Computing at the Hong Kong Polytechnic University. His research interests include

cyber-physical systems, real-time/embedded systems and networking, wireless technology, and their applications in industrial control, medicine, and assisted living. He has published about 30 research papers on various venues. He has received a best paper award from the IEEE Transactions on Industrial Informatics (2008). Dr. Wang is also a member of the ACM.



Lei Bu (M'12) is an assistant professor in the Department of Computer Science and Technology, State Key Laboratory for Novel Software Technology at Nanjing University, P.R.China. He received his B.S. and PH.D. degree in Computer Science from Nanjing University in 2004 and 2010 respectively. He has been visiting scholar in Carnegie Mellon University, Fondazione Bruno Kessler, and University of Texas at Dallas. His main research interests include formal method, model checking, especially verification of hybrid system and cyber-physical system. He has published more than 30 research papers in major peer-reviewed international journals and conference proceedings. Dr. Bu is also a member of the ACM.

and, as a principal investigator, obtained over HK\$25 million grants. Dr. Cao is a senior member of China Computer Federation, a senior member of IEEE, and a member of ACM. He is the Chair of the Technical Committee on Distributed Computing of IEEE Computer Society. Dr. Cao has served as an associate editor and a member of the editorial boards of many international journals, and a chair and member of organizing / program committees for many international conferences. Dr. Cao received the BSc degree in computer science from Nanjing University, Nanjing, China, and the MSc and the Ph.D degrees in computer science from Washington State University, Pullman, WA, USA.



Jian-nong Cao Dr. Cao is currently a chair professor and head of the Department of Computing at Hong Kong Polytechnic University. His research interests include parallel and distributed computing, computer networks, mobile and pervasive computing, fault tolerance, and middleware. He has co-authored 4 books, co-edited 9 books, and published over 300 papers in major international journals and conference proceedings. He has directed and participated in numerous research and development projects

and, as a principal investigator, obtained over HK\$25 million grants. Dr. Cao is a senior member of China Computer Federation, a senior member of IEEE, and a member of ACM. He is the Chair of the Technical Committee on Distributed Computing of IEEE Computer Society. Dr. Cao has served as an associate editor and a member of the editorial boards of many international journals, and a chair and member of organizing / program committees for many international conferences. Dr. Cao received the BSc degree in computer science from Nanjing University, Nanjing, China, and the MSc and the Ph.D degrees in computer science from Washington State University, Pullman, WA, USA.



Xue Liu is an associate professor in the School of Computer Science at McGill University. He received his Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 2006. He received his B.S. degree in Mathematics and M.S. degree in Automatic Control both from Tsinghua University, Beijing, China. He has also worked as the Samuel R. Thompson Associate Professor in the University of Nebraska-Lincoln and HP Labs in Palo Alto, California. His research interests are in computer networks and

communications, smart grid, real-time and embedded systems, cyber-physical systems, data centers, and software reliability. Dr. Liu has published more than 150 research papers in major peer-reviewed international journals and conference proceedings, including the Year 2008 Best Paper Award from IEEE Transactions on Industrial Informatics, and the First Place Best Paper Award of the ACM Conference on Wireless Network Security (WiSec 2011).

From Offline toward Real-Time: A Hybrid Systems Model Checking and CPS Co-Design Approach for Medical Device Plug-and-Play Collaborations [Supplementary File]

Tao Li, *Student Member, IEEE*, Feng Tan, Qixin Wang, *Member, IEEE*, Lei Bu, Jian-nong Cao, *Senior Member, IEEE*, and Xue Liu, *Member, IEEE*,



APPENDIX A DECIDABILITY

This section gives a proof for the decidability of SNZ-LHA system. Though [1] gives a guideline for proof, to our best knowledge, a formal proof is still missing in the literature. We therefore give a formal proof in the following.

The proof involves heavy usage of symbols. Due to space limit, we are not going to explicitly re-define each symbol, instead, readers shall refer to [2] for the definitions of the corresponding symbols.

We first define the following concepts:

Definition 1: We say trajectory τ is from hybrid automaton state (v, σ) to a state space χ , denoted as $(v, \sigma) \rightsquigarrow_{\tau} \chi$, iff $\tau(0, 0) = (v, \sigma)$ and $\tau(h, \delta_h) \in \chi$, where $h = \|\tau\|$ is the number of hops of τ .

For a finite-horizon model checking problem on whether $\sigma \models \varphi_1 \exists \mathcal{U}_{\leq T} \varphi_2$, we can use the well-known SMC-procedure proposed in [2]. To prove the decidability of SNZ-LHA system, we only need to prove SMC-procedure has limited iterations for our case.

Let us first prove the following lemmas.

Lemma 1: In SMC-procedure reachability model checking, if automaton state $(v, \sigma) \in \chi_i \setminus \chi_{i-1}$ (see Section 5.1 of [2] for the definition of χ_i), then $\forall \tau$, if $(v, \sigma) \rightsquigarrow_{\tau} \chi_0$, then $\|\tau\| \geq \lfloor \frac{i}{2} \rfloor$.

Proof: Suppose there is a trajectory τ , s.t. $(v, \sigma) \rightsquigarrow_{\tau} \chi_0$ and $h = \|\tau\| < \lfloor \frac{i}{2} \rfloor$.

- This is the supplementary file.
- T. Li, F. Tan, Q. Wang, and J. Cao are with Dept. of Computing, The Hong Kong Polytechnic Univ., Hong Kong S. A. R. Email: csqwang@comp.polyu.edu.hk
- L. Bu is with State Key Lab for Novel Software Technology, Dept. of Computer Sci. and Tech., Nanjing Univ., Nanjing, 210093, P. R. China. Email: bulei@nju.edu.cn
- X. Liu is with School of Computer Science, McGill Univ., Montreal, Canada. Email: xue.liu@mcgill.ca

Note τ can be denoted as $\tau = (v_0, \delta_0, \rho_0) \rightarrow (v_1, \delta_1, \rho_1) \rightarrow \dots \rightarrow (v_h, \delta_h, \rho_h)$, which consists of h transitions between $(h+1)$ v -trajectories, and $(v_0, \rho_0(0)) = (v, \sigma)$ and $(v_h, \rho_h(\delta_h)) \in \chi_0$.

Since in each iteration of SMC-procedure, all predecessor regions within one v -trajectory or within one transition is included, it hence takes no more than $(2h+1)$ iterations for SMC-procedure to include state (v, σ) . Meanwhile, $h < \lfloor \frac{i}{2} \rfloor \Rightarrow h \leq \lfloor \frac{i}{2} \rfloor - 1 \Rightarrow 2h+1 \leq 2(\lfloor \frac{i}{2} \rfloor - 1) + 1 \leq 2(\frac{i}{2} - 1) + 1 = i - 1$. This means $(v, \sigma) \in \chi_{i-1}$, which contradicts $(v, \sigma) \in \chi_i \setminus \chi_{i-1}$. \square

Lemma 2: Let A be an SNZ-LHA system. Then for any trajectory τ in A , the trajectory duration $\delta_{\tau} \geq \lfloor \frac{\|\tau\|}{\|E\|+1} \rfloor \varepsilon$, where $\|\tau\|$ is the hop length of τ , $\|E\|$ is the number of transitions of A , and ε is defined in the definition of SNZ-LHA system.

Proof: Due to the well-known pigeonhole principle, for every sub-trajectory τ' of τ , if $\|\tau'\| \geq \|E\| + 1$, then τ' must have passed at least one cycle of transitions in A . Therefore, trajectory τ must have passed at least $\lfloor \frac{\|\tau\|}{\|E\|+1} \rfloor$ cycles of transitions in A without temporal overlapping. According to the theorem of SNZ-LHA system decidability, every cycle of transitions takes at least ε seconds to pass, the lemma hence holds. \square

Lemma 3: Suppose the LHA system only consists of one LHA A . Let $I = 2(\lceil \frac{T}{\varepsilon} \rceil + 1)(\|E\| + 1) + 1$, where T is the finite-horizon for finite-horizon reachability model checking, ε is defined in the definition of SNZ-LHA system, and $\|E\|$ is the number of transitions in A , then the SMC-procedure on model checking finite-horizon reachability terminates at the $(I+1)$ th iteration.

Proof: Suppose there is automaton state $(v, \sigma) \in \chi_{I+1} \setminus \chi_I$, then $\forall \tau. (v, \sigma) \rightsquigarrow_{\tau} \chi_0$, we have

$$\|\tau\| \geq \lfloor \frac{I+1}{2} \rfloor \quad (\text{due to Lemma 1}),$$

then trajectory duration

$$\begin{aligned}
\delta_\tau &\geq \left\lfloor \frac{\|\tau\|}{\|E\| + 1} \right\rfloor \varepsilon \quad (\text{due to Lemma 2}) \\
&\geq \left\lfloor \frac{\lfloor \frac{I+1}{2} \rfloor}{\|E\| + 1} \right\rfloor \varepsilon = \left\lfloor \frac{\lfloor \frac{2(\lceil \frac{T}{\varepsilon} \rceil + 1)(\|E\| + 1) + 2}{2} \rfloor}{\|E\| + 1} \right\rfloor \varepsilon \\
&= \left\lfloor \frac{\lfloor (\lceil \frac{T}{\varepsilon} \rceil + 1)(\|E\| + 1) + 1 \rfloor}{\|E\| + 1} \right\rfloor \varepsilon \\
&\geq \left\lfloor \frac{(\lceil \frac{T}{\varepsilon} \rceil + 1)(\|E\| + 1)}{\|E\| + 1} \right\rfloor \varepsilon = \left\lfloor \lceil \frac{T}{\varepsilon} \rceil + 1 \right\rfloor \varepsilon \\
&\geq \left\lfloor \frac{T}{\varepsilon} + 1 \right\rfloor \varepsilon > \frac{T}{\varepsilon} \varepsilon = T.
\end{aligned}$$

This violates the assumption that φ_1 and φ_2 already include the requirement that global time $t \in [0, T]$ for the case of finite-horizon reachability model checking. Hence $\chi_{I+1} \setminus \chi_I = \emptyset$, which means SMC-procedure terminates at the $(I + 1)$ th iteration. \square

With the above definitions and lemmas, we can now prove the decidability of SNZ-LHA system:

Because the LHAs in the system can be combined into a single LHA within polynomial time [2], and the resulted LHA A still satisfies the preconditions of SNZ-LHA system, hence the theorem holds due to Lemma 3. \blacksquare

APPENDIX B DECIDABLE DESIGN PATTERN

To describe these patterns, however, we need the following two concepts.

Definition 2 (Cycle of Transitions): Given a hybrid automaton $A = (\vec{x}, \vec{x}^0, V, v^0, inv, dif, E, act, L, syn)$, a *cycle of transitions* is a sequence of $e_0 v_0 e_1 v_1 \dots e_{k-1} v_{k-1} e_0$, where $e_i \in E$, $v_i \in V$, and $e_i = (v_{((i-1) \bmod k)}, v_i)$ ($i = 0, 1, \dots, k - 1$).

Definition 3 (Minimal Dwelling Time): Given a hybrid automaton $A = (\vec{x}, \vec{x}^0, V, v^0, inv, dif, E, act, L, syn)$, a location $v \in V$ has *minimal dwelling time* of ε iff for any trajectory τ that enters v via a transition, τ must stay in v for at least ε time before being able to leave v via a transition.

With the above concepts, we can describe the following design patterns, assuming \mathcal{S} denotes a set of LHAs.

Definition 4 (ε -Minimal Dwelling Time Pattern): Given a cycle of transitions \mathcal{C} , if there is at least one location v in \mathcal{C} , s.t. v has a minimal dwelling time of ε , then \mathcal{C} complies with ε -Minimal Dwelling Time pattern.

Example 1: The supervisor hybrid automaton (see Fig. 8 of the main file) has a cycle of transitions $\mathcal{C}_1 = \text{“eventSupervisorApprove LaserApproved eventNormalDisapprove LaserDisapproved eventSupervisorApprove”}$. Since all transitions entering location LaserDisapproved sets $t_{disapprove}$ to 0; while all

transitions that leaves LaserDisapproved requires $t_{disapprove} \geq T_{disapprove}^{min}$. We assume $T_{disapprove}^{min}$ is a positive constant, then LaserDisapproved has minimal dwelling time of $\varepsilon = T_{disapprove}^{min} > 0$. As LaserDisapproved is in \mathcal{C}_1 , \mathcal{C}_1 hence complies with ε -Minimal Dwelling Time pattern. \blacksquare

Definition 5 (ε -Alternating Cyber-Value Pattern): Given a set \mathcal{S} of LHAs, and suppose LHA $A \in \mathcal{S}$ has a cycle of transitions $\mathcal{C} = e_0 v_0 e_1 v_1 \dots e_{k-1} v_{k-1} e_0$. If there are two transitions e_i, e_j ($i, j \in \{0, 1, \dots, k - 1\}$) in \mathcal{C} , s.t.

- 1) $(i < j) \vee ((i \neq 0) \wedge (j = 0))$;
- 2) to trigger e_i , a state variable $x_l \in \vec{x}$ must *first* perform a discrete value switch from s to s' (possibly by another automaton in \mathcal{S});
- 3) to trigger e_j , the same x_l must equal s ;
- 4) x_l does not change value within any locations (i.e., it only changes during transitions).
- 5) $s \neq s'$, and all transitions in \mathcal{S} that can switch x_l from s to s' enter target locations with a minimal dwelling time of ε .

then \mathcal{C} complies with ε -Alternating Cyber-Value pattern.

Example 2: In the ventilator hybrid automaton (see Fig. 6 of the main file), there is a cycle of transitions $\mathcal{C}_2 = \text{“eventVentResume PumpOut eventVentToHold PumpIn eventVentHold Hold eventVentResume”}$. Note that to trigger eventVentResume, *LaserApprove* must be first switched from **true** to **false**; and to trigger eventVentHold, *LaserApprove* must equal **true**. *LaserApprove* is a computer logic (i.e., cyber-) variable that does not change in any locations. Plus, all transitions that set *LaserApprove* from **true** to **false** enter the LaserDisapproved location (see Fig. 8 of the main file), which has a minimal dwelling time of $\varepsilon = T_{disapprove}^{min}$, where $T_{disapprove}^{min}$ is a positive constant. This implies \mathcal{C}_2 complies with ε -Alternating Cyber-Value pattern. \blacksquare

Definition 6 (ε -Alternating Physical-Value Pattern):

Given a cycle of transitions \mathcal{C} , if there are two transitions e_i, e_j in \mathcal{C} , s.t.

- 1) to trigger e_i , a state variable $x_l \in \vec{x}$ must equal s ;
- 2) to trigger e_j , the same x_l must equal s' ;
- 3) $s \neq s'$, and x_l represent a physical world parameter, whose value can only change continuously, and there is an upper bound $R > 0$ on $|\dot{x}_l|$, s.t., $\frac{2|s-s'|}{R} \geq \varepsilon$.

then \mathcal{C} complies with ε -Alternating Cyber-Value pattern.

Example 3: In the ventilator hybrid automaton (see Fig. 6 of the main file), there is a cycle of transitions $\mathcal{C}_3 = \text{“eventVentPumpOut PumpOut eventVentPumpIn PumpIn eventVentPumpOut”}$. To trigger eventVentPumpOut, state variable H_{vent} must equal 0.3(m); while to trigger eventVentPumpIn, H_{vent} must equal 0(m). Meanwhile, as H_{vent} represents a physical world parameter: the current height of ventilator cylinder. Its value can only change continuously, and the change rate is bounded by $|\dot{H}_{vent}| = 0.1(\text{m/sec})$. Therefore, to change from 0.3(m) to 0(m) and back to 0.3(m), it takes at least $2|0.3(\text{m}) - 0(\text{m})|/0.1(\text{m/sec}) = 6(\text{sec})$. Therefore \mathcal{C}_3 complies

with ε -Alternating Physical-Value pattern, where we can pick $\varepsilon = 6(\text{sec})$. ■

Proof of Theorem 2 of the Main File: If trajectory τ passes through a transition twice, then it passes through a cycle of transitions \mathcal{C} . Given \mathcal{C} must comply with one of the following patterns:

Case 1: ε -Minimal Dwelling Time pattern, then τ must stay in one location on \mathcal{C} for more than ε .

Case 2: ε -Alternating Cyber-Value pattern, then τ must have changed a state variable x_l 's value from s to s' and back to s , where x_l , s , and s' are described in Definition 5. According to Definition 5, this costs at least ε amount of time.

Case 3: ε -Alternating Physical-Value pattern, then τ must have changed a state variable x_l 's value from s to s' and back to s , where x_l , s , and s' are described in Definition 6. According to Definition 6, since x_l is continuous and $|\dot{x}_l| \leq R$, altering x_l from s to s' and back to s takes at least $\frac{2|s-s'|}{R} \geq \varepsilon$ amount of time.

In summary, there must be $\delta_\tau \geq \varepsilon$. Therefore, \mathcal{S} is SNZ-LHA-System, and hence is decidable for finite-horizon reachability model checking. ■

APPENDIX C NP-HARDNESS OF SNZ-LHA SYSTEM

We can reduce the well-known NP-Hard problem of Directed Hamiltonian Cycle to the problem of finite-horizon reachability model checking of an SNZ-LHA-System.

The Directed Hamiltonian Cycle problem is as follows:

Given a directed graph $G = (V_G, E_G)$, where V_G and E_G are its vertex set and directed edge set respectively, does G contain a directed Hamiltonian cycle?

Given an instance of Directed Hamiltonian Cycle problem P_G on directed graph $G = (V_G, E_G)$, we can construct an SNZ-LHA-System $A = (\vec{x}, \vec{x}^0, V, v^0, \text{inv}, \text{dif}, E, \text{act}, L, \text{syn})$ in polynomial time, where

$V = V_G$, i.e., each vertex in G is regarded as a location in A . Denote $n = |V| = |V_G|$, and hence denote $V = V_G = \{v_i\}$, where $i = 1, 2, \dots, n$.

$E = E_G$, i.e., each edge in G is regarded as an edge in A .

$\vec{x} = (x_1, x_2, \dots, x_n, t)$, i.e., we assign a variable x_i for each $v_i \in V = V_G$ (where $i = 1, 2, \dots, n$); and t is a timer for fixed dwelling time guarantee, which will be explained later.

$\vec{x}^0 = (0, 0, \dots, 0)$, i.e., x_i (where $i = 1, 2, \dots, n$) and t are all initialized to 0.

v^0 can be any location. Without loss of generality, let us pick $v^0 = v_1$.

For each $v_i \in V$ (where $i = 1, 2, \dots, n$), $\text{inv}(v_i)$ corresponds to the inequality proposition of

$$0 \leq t < 1. \quad (1)$$

For each $v_i \in V$ (where $i = 1, 2, \dots, n$), $\text{dif}(v_i)$ corre-

sponds to the following continuous activities:

$$\begin{aligned} \dot{x}_i &= 1, \\ \dot{x}_k &= 0, \text{ for each } k \neq i \text{ and } k \in \{1, 2, \dots, n\} \\ \dot{t} &= 1. \end{aligned} \quad (2)$$

For each edge $e = (v_i, v_j) \in E$ (where $i, j \in \{1, 2, \dots, n\}$), $\text{act}(e)$ corresponds to the following discrete actions:

$$t = 1, \quad (3)$$

$$t' = 0. \quad (4)$$

Formulae (3), (4), (1), and (2) together imply every location $v \in V$ has a *fixed* dwelling time of 1: once entered, one has to stay exactly 1 unit of time, and then move to another location. This further implies LHA A complies with the ε -Minimal Dwelling Time Pattern, where $\varepsilon = 1$. According to Theorem of decidable design pattern, LHA A henceforth is a SNZ-LHA-System.

Finally, we choose $L = \emptyset$ and $\text{syn} = \emptyset$ as they are irrelevant to our proof.

With SNZ-LHA-System A at hand, our finite-horizon SNZ-LHA-System reachability model checking problem P_A checks whether the following state space S is reachable within finite-horizon $T = n + 1$:

$$\begin{aligned} S &= \{(v_1, \vec{x}) \mid \text{where } 1.5 < x_1 < 2, \\ &\text{and } 0.5 < x_i \leq 1 \text{ for each } i = 2, 3, \dots, n\}. \end{aligned}$$

A “yes” answer to P_A implies there is a cyclic trajectory on A that traverse each vertex $v \in V$ exactly once and returns to the initial location of v_1 . This trajectory hence corresponds to a Hamiltonian Cycle in G , hence a “yes” answer to P_G . Conversely, a “yes” answer to P_G implies there is a Hamiltonian Cycle in G , along this cycle, we can traverse A to reach S , hence a “yes” answer to P_A .

From above, we prove Directed Hamiltonian Cycle problem can be reduced to finite-horizon reachability model checking of an SNZ-LHA-System problem in polynomial time. As Directed Hamiltonian Cycle problem is NP-Hard, finite-horizon reachability model checking of an SNZ-LHA-System problem is hence NP-Hard. ■

APPENDIX D FALSE NEGATIVES AND FALSE POSITIVES DETAILED DISCUSSIONS

For ease of narration, we call our proposed online model checking based MDPnP practice as “*MDPNP-practice*”; call the corresponding online modeling and online model checking as “*MDPNP-online-modeling*” and “*MDPNP-online-model-checking*” respectively.

Still take the laser tracheotomy for example, Table 1 lists all possible cases for “*MDPNP-practice*”. We see that the upper bounds of accident probability are

$$P_m^{\text{cons}} = p(+)\text{p}_m(-|+) \quad (5)$$

$$\text{and } P_m^{\text{aggr}} = p(+)[\text{p}_m(-|+) + \text{p}_m(?|+)] \quad (6)$$

respectively for “conservative mode” and “aggressive mode”, where $p(+)$ is the probability to reach unsafe states under the

absolutely accurate patient model (i.e. the patient model in “God’s view”, rather than the model described by Fig. 4 of the main file); $p_m(-|+)$, $p_m(?|+)$ are respectively the conditional probability that MDPnP-online-model-checking gives “negative” answer (i.e., a false-negative answer), or misses deadline (i.e., cannot give a deterministic answer). Note false-positive is not a big concern as it will trigger fall-back plan, leaving no chances for accidents (though may be a nuisance to the surgeon).

TABLE 1
All Possible Cases for MDPnP Practice

| Reality | Online-Model-Checking Result | What Happens | Accident Possible? |
|----------|------------------------------|--|--------------------|
| positive | positive | scenario 1 | No |
| | negative | scenario 2 | Yes |
| | deadline miss | scenario 1 (cons) scenario 2 (aggr) | No Yes |
| negative | no need to care | no need to care | No |

positive: unsafe states are reachable.

negative: unsafe states are not reachable.

scenario 1: fall-back plan kicks in, which forbids use of laser and keeps ventilator on; the worst case is that the surgeon may be annoyed.

scenario 2: the system run as what Fig. 4, 6, 7, 8 of the main file describe.

In comparison to MDPnP-practice, now let us study the *current-practice* (i.e., the actual practice in nowadays hospitals) of laser tracheotomy.

First, the role of supervisor (i.e. the procedure described in Fig. 8 of the main file) is taken over by a human-supervisor. Usually, the human-supervisor is the surgeon himself/herself; but for clarity, let us differentiate the two persons.

Second, as for line 4 of the algorithm described in Fig. 10 of the main file, instead of MDPnP-online-modeling, the human-supervisor uses his/her subjective judgement to model the patient in the near future (e.g., replace $Sp\dot{O}_2(t_0)$ in Fig. 4 of the main file with his/her subjective prediction). We call this “*subjective-online-modeling*”.

Third, as for line 6 and 11 of the algorithm described in Fig. 10 of the main file, instead of MDPnP-online-model-checking, the human-supervisor uses his subjective judgement to decide whether unsafe states are reachable. We call this “*subjective-online-model-checking*”.

Therefore, reusing the same analysis on the MDPnP-practice, we can derive the upper bounds of accident probability for the current-practice:

$$P_c^{cons} = p(+)p_s(-|+) \quad (7)$$

$$\text{and } P_c^{aggr} = p(+) [p_s(-|+) + p_s(?|+)] \quad (8)$$

respectively for “conservative mode” and “aggressive mode”, where $p_s(-|+)$, $p_s(?|+)$ are respectively the conditional probability that subjective-online-model-checking gives “negative” answer (i.e., a false-negative answer), or misses deadline (i.e., cannot give a deterministic answer).

Suppose we adopt the “conservative mode”. By comparing Equation (5) and (7), we see the MDPnP-practice is safer than the current-practice when

$$p_m(-|+) \leq p_s(-|+). \quad (9)$$

How to mathematically verify Inequality (9) is beyond the scope of this paper. However, we can still verify empirically. For example, if some well-established math model for predicting patient near-future behavior exists [3], then we’d better use MDPnP-online-modeling rather than relying on subjective-online-modeling. Or, we can carry out comparison using well-known benchmark patient traces, to see which online-modeling is more trustworthy.

The same thing is for “aggressive mode”, except that Inequality (9) now becomes

$$p_m(-|+) + p_m(?|+) \leq p_s(-|+) + p_s(?|+). \quad (10)$$

APPENDIX E WIRELESS COMMUNICATIONS LINKS DETAILED DISCUSSIONS

To evaluate the method proposed in Section 6.2 of the main file, we redo the evaluations of Section 5.1 of the main file to redraw its Table 1 and 2. All other settings are the same except that we consider the link between the oximeter and supervisor to be a wireless link.

The wireless link is unreliable, therefore packets carrying blood oxygen level readings may be lost. The packet losses are treated with the method proposed in Section 6.2 of the main file. The results are summarized in Table 2 and 3 of this supplementary file. These two tables respectively replace Table 1 and 2 of Section 5.1 of the main file.

The results show that with packet loss rate of $\leq 3\%$, both the deadline miss rates and the relative errors are moderately low. Hence even with the unreliable wireless link, we can still carry out online hybrid model checking.

Note with the new advancements in medical grade wireless communications technology, it is possible to control wireless link packet loss rate to below 1%, or even 0.1% [4][5].

TABLE 2

Statistics of execution time cost of online model checking (unit: second; deadline $D = 2$ seconds)

| | Packet loss rate | % of trials missed deadline | Execution time of those caught deadline (secs) | | | |
|-----------------|------------------|-----------------------------|--|-------|-------|-------|
| | | | Min | Max | Mean | Std |
| PhysioNet Trace | 1% | 5.034% | 0.817 | 1.720 | 0.932 | 0.124 |
| | 2% | 6.387% | 0.817 | 1.720 | 0.932 | 0.125 |
| | 3% | 11.429% | 0.817 | 1.720 | 0.929 | 0.124 |
| HKPolyU Trace | 1% | 4.690% | 0.818 | 1.940 | 0.964 | 0.147 |
| | 2% | 6.030% | 0.818 | 1.940 | 0.964 | 0.146 |
| | 3% | 11.725% | 0.818 | 1.940 | 0.966 | 0.149 |

TABLE 3

Statistics of blood oxygen level online modeling relative errors (%)

| | Packet loss rate | Min | Max | Mean | Std |
|-----------------|------------------|--------|-------|-------|-------|
| PhysioNet Trace | 1% | < 0.01 | 2.529 | 0.512 | 0.526 |
| | 2% | < 0.01 | 2.529 | 0.510 | 0.523 |
| | 3% | < 0.01 | 2.529 | 0.503 | 0.526 |
| HKPolyU Trace | 1% | < 0.01 | 3.918 | 0.602 | 0.602 |
| | 2% | < 0.01 | 3.918 | 0.600 | 0.606 |
| | 3% | < 0.01 | 3.918 | 0.591 | 0.608 |

REFERENCES

- [1] O. Roux and V. Rusu, "Uniformity for the decidability of hybrid automata," *Static Analysis*, vol. 1145, pp. 301–316, 1996.
- [2] R. Alur *et al.*, "Automatic symbolic verification of embedded systems," *IEEE Trans. on Software Engineering*, vol. 22, no. 3, pp. 181–201, Mar. 1996.
- [3] J. X. Mazoit, K. Butscher, and K. Samii, "Morphine in postoperative patients: Pharmacokinetics and pharmacodynamics of metabolites," *Anesthesia and Analgesia*, vol. 105, no. 1, pp. 70–78, 2007.
- [4] J. Huang *et al.*, "Beyond co-existence: Exploiting WiFi white space for ZigBee performance assurance," *Proc. of ICNP*, pp. 305–314, Oct. 2010.
- [5] Q. Wang *et al.*, "Building robust wireless LAN for industrial control with the DSSS-CDMA cell phone network paradigm," *IEEE Trans. on Mobile Computing*, vol. 6, no. 6, pp. 706–719, Jun. 2007.