

# Decentralized Real-Time Safety Verification for Distributed Cyber-Physical Systems

Hoang-Dung Tran<sup>1</sup>, Luan Viet Nguyen<sup>2</sup>, Patrick Musau<sup>1</sup>, Weiming Xiang<sup>1</sup>,  
and Taylor T. Johnson<sup>1</sup>

<sup>1</sup> Institute for Software Integrated Systems, Vanderbilt University, TN, USA

<sup>2</sup> Department of Computer and Information Science, University of Pennsylvania, PA,  
USA

**Abstract.** Safety-critical distributed cyber-physical systems (CPSs) have been found in a wide range of applications. Notably, they have displayed a great deal of utility in intelligent transportation, where autonomous vehicles communicate and cooperate with each other via a high-speed communication network. Such systems require an ability to identify maneuvers in real-time that cause dangerous circumstances and ensure the implementation always meets safety-critical requirements. In this paper, we propose a real-time decentralized safety verification approach for a distributed multi-agent CPS with the underlying assumption that all agents are time-synchronized with a low degree of error. In the proposed approach, each agent periodically computes its local reachable set and exchanges this reachable set with the other agents with the goal of verifying the system safety. Our method, implemented in Java, takes advantages of the timing information and the reachable set information that are available in the exchanged messages to reason about the safety of the whole system in a decentralized manner. Any particular agent can also perform local safety verification tasks based on their local clocks by analyzing the messages it receives. We applied the proposed method to verify, in real-time, the safety properties of a group of quadcopters performing a distributed search mission.

## 1 Introduction

The emergence of 5G technology has inspired a massive wave of the research and development in science and technology in the era of IoT where the communication between computing devices has become significantly faster with lower latency and power consumption. The power of this modern communication technology influences and benefits all aspects of Cyber-Physical Systems (CPSs) such as smart grids, smart homes, intelligent transportation and smart cities. In particular, the study of autonomous vehicles has become an increasingly popular research field in both academic and industrial transportation applications. Automotive crashes pose significant financial and life-threatening risks, and there is an urgent need for advanced and scalable methods that can efficiently verify a distributed system of autonomous vehicles.

Over the last two decades, although many methods have been developed to conduct reachability analysis and safety verification of CPS, such as the approaches proposed in [1, 4, 10, 11, 13, 15, 18], applying these techniques to *real-time distributed* CPS remains a big challenge. This is due to the fact that, 1) all existing techniques have intensive computation costs and are usually too slow to be used in a real-time manner and, 2) these techniques target the safety verification of a *single* CPS, and therefore they naturally cannot be applied efficiently to a *distributed* CPS where clock mismatches and communication between agents (i.e., individual systems) are essential concerns. Since the future autonomous vehicles systems will work distributively involving effective communication between each agent, there is an urgent need for an approach that can provide formal guarantees of the safety of distributed CPS in real-time. More importantly, the safety information should be defined based on the *agents local clocks* to allow these agents to perform “intelligent actions” to escape from the upcoming dangerous circumstances. For example, if an agent A knows based on its local clock that it will collide with an agent B in the next 5 seconds, it should perform an action such as stopping or quickly finding a safe path to avoid the collision.

In this paper, we propose a *decentralized real-time safety verification* approach for a distributed CPS with multiple agents. We are particularly interested in two types of safety properties. The first one is a *local safety property* which specifies the local constraints of the agent operation. For example, each agent is only allowed to move within a specific region, does not hit any obstacles, and its velocity needs to be limited to specific range. This type of property does not require the information of other agents and can be verified locally at run-time. The second safety property is a *global safety property* in which we want to check if there are any potential collision occurring between the agents.

Our decentralized real-time safety verification approach works as follows. Each agent *locally* and *periodically computes* the local reachable set from the current local time to the next  $T$  seconds, and then *encodes* and *broadcasts* its reachable set information to the others via a communication network. When the agent receives a reachable set message, it immediately *decodes* the message to read the reachable set information of the sender, and then performs *peer-to-peer collision checking* based on its current state and the reachable set of the sender. Additionally, the local safety property of the agent is verified simultaneously with the reachable set computation process at run-time. The proposed verification approach is based on an underlying assumption that is, all agents are time-synchronized to some level of accuracy. This assumption is reasonable as it can be achieved by using existing time synchronization protocols such as the Network Time Protocol (NTP). Our approach has successfully verified in real-time the local safety properties and collision occurrences for a group of quadcopters conducting a search mission.

## 2 Problem Formulation

In this paper, we consider a distributed CPS with  $N$  agents that can communicate with each other via an asynchronous communication channel.

*Communication Model* The communication between agents is implemented by the *actions* of sending and receiving messages over an asynchronous communication channel. We formally model this communication model as a single automaton, **Channel**, which stores the set of in-flight messages that have been sent, but are yet to be delivered. When an agent sends a message  $m$ , it invokes a  $send(m)$  action. This action adds  $m$  to the *in-flight* set. At any arbitrary time, the **Channel** chooses a message in the in-flight set to either delivers it to its recipient or removes it from the set. All messages are assumed to be unique and each message contains its sender and recipient identities. Let  $M$  be the set of all possible messages used in communication between agents. The sending and receiving messages by agent  $i$  are denoted by  $M_{i,*}$  and  $M_{*,i}$ , respectively.

*Agent Model* The  $i^{th}$  agent is modeled as a hybrid automaton [12,22] defined by the tuple  $\langle \mathcal{A}_i = V_i, A_i, \mathcal{D}_i, \mathcal{T}_i \rangle$ , where:

- a)  $V_i$  is a set of variables consisting of the following: i) a set of continuous variables  $X_i$  including a special variable  $clk_i$  which records the agent's *local time*, and ii) a set of discrete variables  $Y_i$  including the special variable  $msghist_i$  that records all sent and received messages. A valuation  $\mathbf{v}_i$  is a function that associates each  $v_i \in V_i$  to a value in its type. We write  $val(V_i)$  for the set of all possible valuations of  $V_i$ . We abuse the notion of  $\mathbf{v}_i$  to denote a state of  $\mathcal{A}_i$ , which is a valuation of all variables in  $V_i$ . The set  $Q_i \triangleq val(V_i)$  is called the set of *states*.
- b)  $A_i$  is a set of *actions* consisting of the following subsets: i) a set  $\{send_i(m) \mid m \in M_{i,*}\}$  of send actions (i.e., output actions), ii) a set  $\{receive_i(m) \mid m \in M_{*,i}\}$  of receive actions (i.e., input actions), and iii) a set  $H_i$  of other, ordinary actions.
- c)  $\mathcal{D}_i \subseteq val(V_i) \times A_i \times val(V_i)$  is called the set of *transitions*. For a transition  $(\mathbf{v}_i, a_i, \mathbf{v}'_i) \in \mathcal{D}_i$ , we write  $\mathbf{v}_i \xrightarrow{a_i} \mathbf{v}'_i$  in short. i) If  $a_i = send_i(m)$  or  $receive_i(m)$ , then all the components of  $\mathbf{v}_i$  and  $\mathbf{v}'_i$  are identical except that  $m$  is added to  $msghist$  in  $\mathbf{v}'_i$ . That is, the agent's other states remain the same on message sends and receives. Furthermore, for every state  $\mathbf{v}_i$  and every receive action  $a_i$ , there must exist a  $\mathbf{v}'_i$  such that  $\mathbf{v}_i \xrightarrow{a_i} \mathbf{v}'_i$ , i.e., the automaton must have well-defined behavior for receiving any message in any state. ii) If  $a_i \in H_i$ , then  $\mathbf{v}_i.msghist = \mathbf{v}'_i.msghist$ .
- d)  $\mathcal{T}_i$  is a collection of trajectories for  $X_i$ . Each trajectory of  $X_i$  is a function mapping an interval of time  $[0, t], t \geq 0$  to  $val(V_i)$ , following a flow rate that specifies how a real variable  $x_i \in X_i$  evolving over time. We denote the *duration* of a trajectory as  $\tau_{dur}$ , which is the right end-point of the interval  $t$ .

*Agent Semantics* The *behavior* of each agent can be defined based on the concept of an *execution* which is a particular run of the agent. Given an initial state  $\mathbf{v}_i^0$ , an *execution*  $\alpha_i$  of an agent  $A_i$  is a sequence of states starting from  $\mathbf{v}_i^0$ , defined as  $\alpha_i = \mathbf{v}_i^0, \mathbf{v}_i^1, \dots$ , and for each index  $j$  in the sequence, the state update from  $\mathbf{v}_i^j$  to  $\mathbf{v}_i^{j+1}$  is either a transition or trajectory. A state  $\mathbf{v}_i^j$  is *reachable* if there exists an executing that ends in  $\mathbf{v}_i^j$ . We denote  $\text{Reach}(A_i)$  as the reachable set of agent  $A_i$ .

*System Model* The formal model of the complete system, denoted as **System**, is a network of hybrid automata that is obtained by parallel composing the agent's models and the communication channel. Formally, we can write,  $\text{System} \triangleq \mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_N \parallel \text{Channel}$ . Informally, the agent  $\mathcal{A}_i$  and the communication channel **Channel** are synchronized through sending and receiving actions. When the agent  $A_i$  sends a message  $m \in M_{i,j}$  to the agent  $A_j$ , it triggers the  $\text{send}_i(m)$  action. At the same time, this action is synchronized in the **Channel** automaton by putting the message  $m$  in the *in-flight* set. After that, the **Channel** will trigger (non-deterministically) the  $\text{receive}_j(m)$  action. This action is synchronized in the agent  $A_j$  by putting the message  $m$  into the *msghist* <sub>$j$</sub> .

In this paper, we investigate two real-time safety verification problems for distributed cyber-physical systems as defined in the following.

*Problem 1 (Local safety verification in real-time).* The real-time local safety verification problem is to compute online the reachable set  $\text{Reach}(A_i)$  of the agent and verify if it violates the local safety property, i.e., checking  $\text{Reach}(A_i) \cap \mathcal{U}_i = \emptyset?$ , where  $\mathcal{U}_i \triangleq C_i x_i \leq d_i, x_i \in X_i$  is the unsafe set of the agent.

*Problem 2 (Decentralized real-time collision verification).* The decentralized real-time collision verification problem is to reason in real-time whether an agent  $A_i$  will collide with other agents from its current local time  $t_c^i$  to the *computable, safe time instance in the future*  $T_{safe}$  based on i) the *clock mismatches*, and ii) the *exchanging reachable set messages* between agents. Formally, we require that  $\forall t_c^i \leq t \leq T_{safe}, d_{ij}(t) \geq l$ , where  $d_{ij}(t)$  is the distance between agents  $A_i$  and  $A_j$  at the time  $t$  of the agent  $A_i$  local clock, and  $l$  is the allowable safe distance between agents.

### 3 Real-Time Local Safety Verification

The first important step in our approach is, each agent  $A_i$  computes forwardly its reachable set of states from the current local time  $t^i$  to the next  $(t^i + T)$  seconds which is defined by  $\mathcal{R}_i[t^i, t^i + T]$ . Since there are many variables used in the agent modeling that are irrelevant in safety verification, we only need to compute the reachable set of state that is related to the agent's physical dynamics (so called as *motion dynamics*) which is defined by a nonlinear ODE  $\dot{x}_i = f(x_i, u_i)$ , where  $x_i \in \mathbb{R}^n$  is state vector and  $u_i \in \mathbb{R}^m$  is the control input vector. The agent can switch from one mode to the another mode via discrete transitions, and in each

mode, the control law may be different. When the agent computes its reachable set, the only information it needs are its current set of states  $x_i(t^i)$  and the current control input  $u_i(t^i)$ . It should be clarified that although the control law may be different among modes, the control signal  $u_i$  is updated with the same control period  $T_c^i$ . Consequently,  $u_i$  is a constant vector in each control period.

Assuming that the agent's current time is  $t_j^i = j \times T_c$ , using its local sensors and GPS, we have the current state of the agent  $x_i$ . Note that the local sensors and the provided GPS can only provide the information of interest to some accuracy, therefore the actual state of the agent is in a set  $x_i \in I_i$ . The control signal  $u_i$  is computed based on the state  $x_i$  and a reference signal, e.g., a set point denoting where the agent needs to go to, and then computed control signal is applied to the actuator to control the motion of the agent. From the current set of states  $I_i$  and the control signal  $u_i$ , we can compute the forward reachable set of the agent for the next  $t_j^i + T$  seconds. This reachable set computation needs to be completed after an amount of time  $T_{runtime}^i < T_c^i$  because if  $T_{runtime}^i \geq T_c^i$ , a new  $u_i$  will be updated. The control period  $T_c^i$  is chosen based on the agent's motion dynamics, and thus to control an agent with fast dynamics, the control period  $T_c^i$  needs to be sufficiently small. This is the source of the requirement that the allowable run-time for reachable set computation be small.

To compute the reachable set of an agent in real-time, we use the well-known face-lifting method [3, 6] and a *hyper-rectangle* to represent the reachable set. This method is useful for short-time reachability analysis of real-time systems. It allows users to define an allowable run-time  $T_{runtime}^i$ , and has no dynamic data structures, recursion, and does not depend on complex external libraries as in other reachability analysis methods. More importantly, the accuracy of the reachable set computation can be iteratively improved based on the *remaining allowable run-time*.

Algorithm 3.1 describes the real-time reachability analysis for one agent. The Algorithm works as follows. The time period  $[t^i, t^i + T]$  is divided by  $M$  steps. The reach time step is defined by  $h_i = T/M$ . Using the reach time step and the current set  $I_i$ , the face-lifting method performs a *single-face-lifting operation*. The results of this step are a new reachable set and a *remaining reach time*  $T_{remainReachTime}^i < T$ . This step is iteratively called until the reachable set for the whole time period of interest  $[t^i, t^i + T]$  is constructed completely, i.e., the remaining reach time is equal to zero. Interestingly, with the reach time step size  $h_i$  defined above, the face-lifting algorithm may be finished quickly after an amount of time which is smaller than the allowable run-time  $T_{runtime}^i$  specified by user, i.e., there is still an amount of time called remaining run time  $T_{remainRunTime}^i < T_{runtime}^i$  that is available for us to recall the face-lifting algorithm with a smaller reach time step size, for example, we can recall the face-lifting algorithm with a new reach time step  $h_i/2$ . By doing this, the conservativeness of the reachable set can be iteratively improved. The core step of face-lifting method is the single-face-lifting operation. We refer the readers to [3] for further detail. As mentioned earlier, the local safety property of each agent can be verified at run-time simultaneously with the reachable set computation

---

**Algorithm 3.1** Real-time reachability analysis for agent  $A_i$ .

---

**Input:**  $I_i, u_i, t^i, T, h_i, T_{runtime}^i, \mathcal{U}_i$   
**Output:**  $\mathcal{R}_i[t^i, t^i + T]$ ,  $safe = true$  or  $safe = uncertain$

```
1: procedure INITIALIZATION
2:    $step = h_i$            % Reach time step
3:    $T_1^i = T_{runtime}^i$    % Remaining run-time
4: procedure REACHABILITY ANALYSIS
5:   while ( $T_1^i > 0$ ) do
6:      $\mathcal{CR} = I_i$        % Current reachable set
7:      $safe = true$ 
8:      $T_2^i = T$          % Remaining reach time
9:     while  $T_2^i > 0$  do
10:      % Do Single Face Lifting
11:       $\mathcal{R}, T' = SFL(\mathcal{CR}, step, T_2^i, u_i)$ 
12:       $\mathcal{CR} = \mathcal{R}$      % Update reach set
13:       $T_2^i = T'$      % Update remaining reach time
14:      if ( $\mathcal{CR} \cap \mathcal{U}_i \neq \emptyset$ ) then:  $safe = uncertain$ 
15:       $\mathcal{R}_i[t^i, t^i + T] = \mathcal{CR}$ 
16:      % Update remaining runtime
17:       $T_1^i = T_1^i - (A_i.currentTime() - t^i)$ 
18:      if  $T_1^i \leq 0$  then break
19:      else
20:         $step = h_i/2$    % Reduce reach time step
21:      return  $\mathcal{R}_i[t^i, t^i + T] = \mathcal{CR}, safe$ 
```

---

process. Precisely, let  $\mathcal{U}_i \triangleq C_i x_i \leq d_i$  be the unsafe region of the  $i^{th}$  agent, the agent is said to be safe from  $t^i$  to  $t^i + t \leq t^i + T$  if  $\mathcal{R}_i[t^i, t^i + t] \cap \mathcal{U}_i = \emptyset$ . Since the reachable set  $\mathcal{R}_i[t^i, t^i + t]$  is given by the face-lifting method at run-time, the local safety verification problem for each agent can be solved at run-time. Since the Algorithm 3.1 computes an over-approximation of the reachable set of each agent in a short time interval, it guarantees the soundness of the result as described in the following lemma.

**Lemma 1.** [3, 6] *The real-time reachability analysis algorithm is sound, i.e., the computed reachable set contains all possible trajectories of agent  $A_i$  from  $t^i$  to  $t^i + T$ .*

## 4 Decentralized Real-Time Collision Verification

Our collision verification scheme is performed based on the exchanged reachable set messages between agents. For every control period  $T_c$ , each agent executes the real-time reachability analysis algorithm to check if it is locally safe and to obtain its current reachable set with respect to its current control input. When the current reachable set is available, the agent encodes the reachable set in a

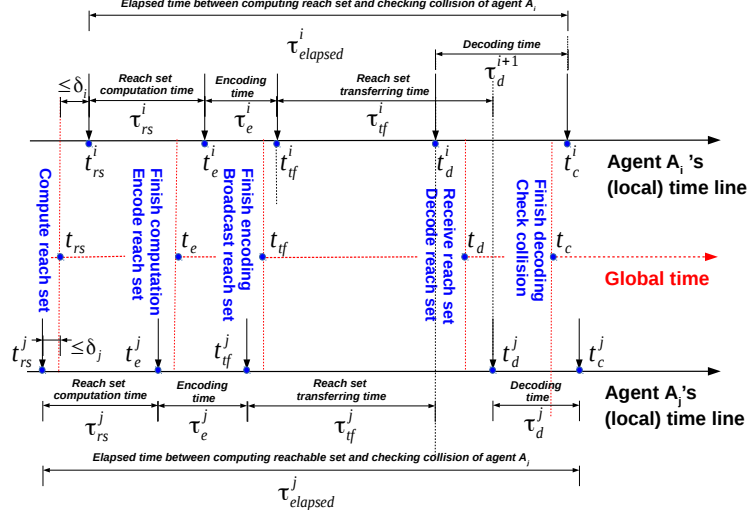


Fig. 1: Timeline for reachable set computing, encoding, transferring, decoding and collision checking.

message and then broadcasts this message to its cooperative agents and listens to the upcoming messages sent from these agents. When a reachable set message arrives, the agent immediately decodes the message to construct the current reachable set of the sender and then performs peer-to-peer collision detection. The process of computing, encoding, transferring, decoding of the reachable set along with collision checking is illustrated in Figure 1 based on the agent's local clock.

Let  $t_{rs}^i$ ,  $t_e^i$ ,  $t_{tf}^i$ ,  $t_d^i$ , and  $t_c^i$  respectively be the instants that we compute, encode, transfer, decode the reachable set and do collision checking on the agent  $A_i$ . Note that these time instants are based on the agent  $A_i$ 's local clock. The actual run-times are defined as follows.

$$\begin{aligned}
 \tau_{rs}^i &= t_e^i - t_{rs}^i, \% \text{ reachablet set computation time,} \\
 \tau_e^i &= t_{tf}^i - t_e^i, \% \text{ encoding time,} \\
 \tau_{tf}^i &\approx t_d^j - t_{tf}^i, \% \text{ transferring time,} \\
 \tau_d^i &= t_c^i - t_d^i, \% \text{ decoding time.}
 \end{aligned}$$

Note that we do not know the exact transfer time  $\tau_{tf}^i$  since it depends on two different local time clocks. The above transfer time formula describes its approximate value when neglecting the mismatch between the two local clocks. The actual reachable set computation time is close to the allowable run-time chosen by user, i.e.,  $\tau_{rs}^i \approx T_{runtime}^i$ . We will see later that the encoding time and decoding time are fairly small in comparison with the transferring time, i.e.,  $\tau_e^i \approx \tau_d^i \ll \tau_{tf}^i$ . All of these run-times provide useful information for selecting an appropriate control period  $T_c$  for an agent. However, for collision checking pur-




Fig. 2: Useful reachable set.

pose, we only need to consider the time instants that an agent starts computing reachable set  $t_{rs}^i$  and checking collision  $t_c^i$ .

A reachable set message contains three pieces of information: the reachable set which is a list of intervals, the time period (based on the local clock) in which this reachable set is valid, i.e., the start time  $t_{rs}^i$  and the end time  $t_{rs}^i + T$  and the time instant that this message is sent. Based on the timing information of the reachable set and the time-synchronization errors, an agent can examine whether or not a received reachable set contains information about the future behavior of the sent agent which is useful for checking collision. The usefulness of the reachable sets used in collision checking is defined as follows.

**Definition 1 (Useful reachable sets).** Let  $\delta_i$  and  $\delta_j$  respectively be the time-synchronization errors of agent  $A_i$  and  $A_j$  in comparison with the virtual global time  $t$ , i.e.,  $t - \delta_i \leq t^i \leq t + \delta_i$  and  $t - \delta_j \leq t^j \leq t + \delta_j$ , where  $t^i$  and  $t^j$  are current local times of  $A_i$  and  $A_j$  respectively. The reachable sets  $\mathcal{R}_i[t_{rs}^i, t_{rs}^i + T]$  and  $\mathcal{R}_j[t_{rs}^j, t_{rs}^j + T]$  of the agent  $A_j$  that are available at the agent  $A_i$  at time  $t_c^i$  are useful for checking collision between  $A_i$  and  $A_j$  if:

$$\begin{aligned} t_c^i &< t_{rs}^j + T - \delta_i - \delta_j, \\ t_c^i &< t_{rs}^i + T. \end{aligned} \quad (1)$$

Assume that we are at a time instant where the agent  $A_i$  checks if a collision occurs. This means that the current local time is  $t_c^i$ . Note that agent  $A_i$  and  $A_j$  are synchronized to the global time with errors  $\delta_i$  and  $\delta_j$  respectively. The reachable set  $\mathcal{R}_j[t_{rs}^j, t_{rs}^j + T]$  is useful if it contains information about the *future behavior* of agent  $A_j$  under the view of the agent  $A_i$  based on its local clock. This can be guaranteed if we have:  $t_{rs}^j + T \geq t_{rs}^i - \delta_j + T > t_c^i + \delta_i$ . Additionally, the current reachabset of agent  $A_i$  contains information about its future behavior if  $t_c^i < t_{rs}^i + T$  as depicted in Figure 2. We can see that if  $t_c^i > t_{rs}^j + T + \delta_i + \delta_j$ , then the reachable set of  $A_j$  contains a past information, and thus it is useless for checking collision. One interesting case is when  $t_{rs}^j + T - \delta_i - \delta_j < t_c^i < t_{rs}^j + T + \delta_i + \delta_j$ . In this case, we do not know whether the received reachable set is useful or not.

*Remark 1.* We note that the proposed approach does not rely on the concept of Lamport happens-before relation [17] to compute the local reachable set of



---

**Algorithm 4.2** Decentralized Real-Time Collision Verification at Agent  $A_i$ .

---

**Input:**  $l$ , % safe distance between agents

**Output:**  $\text{collision}, T_{\text{safe}}$  % collision flag and safe time interval in the future

```
1: procedure PEER-TO-PEER COLLISION DETECTION
2:   if new message  $\mathcal{R}_j[t_{rs}^j, t_{rs}^j + T]$  arrive then
3:     decode message
4:      $t_c^i = A_i.\text{current.time}()$  % current time
5:      $t_{rs}^i = \mathcal{R}_i.t_{rs}^i$  % current reachable set start time
6:     if  $t_c^i < t_{rs}^j + T - \delta_i - \delta_j$  and  $t_c^i < t_{rs}^i + T$  then % check usefulness
7:       compute possible minimum distance  $d_{min}$  between two agents
8:       if  $d_{min} > l$  then
9:         Collision = false
10:         $T_{\text{safe}} = \min(t_{rs}^j + T - \delta_i - \delta_j, t_{rs}^i + T)$ 
11:       else
12:         Collision = uncertain,  $T_{\text{safe}} = []$ 
13:       store the message
```

---

each agent. If the agent could not receive reachable messages from others until a requested time-stamp expires, it still calculates the local reachable set based on its current state and the state information of other agents in the messages it received previously. In other words, our method does not require the reachable set of each agent to be computed corresponding to the ordering of the events (sending or receiving a message) in the system, but only relies on the local clock period and the time-synchronization errors between agents. Such implementation ensures that the computation process can be accomplished in real-time, and is not affected by the message transmission delay.

The peer-to-peer collision checking procedure depicted in Algorithm 4.2 works as follows: when a new reachable set message arrives, the receiving agent decodes the message and checks the usefulness of the received reachable set and its current reachable set. Then, the agent combines its current reachable set and the received reachable set to compute the minimum possible distance between two agents. If the distance is larger than an allowable threshold  $l$ , there is no collision between two agents in some known time interval in the future, i.e.,  $T_{\text{safe}}$ .

**Lemma 2.** *The decentralized real-time collision verification algorithm is sound.*

*Proof.* From Lemma 1, we know that the received reachable set  $\mathcal{R}_j[t_{rs}^j, t_{rs}^j + T]$  contains all possible trajectories of the agent  $A_j$  from  $t_{rs}^j$  to  $t_{rs}^j + T$ . Also, the current reachable set of the agent  $A_i$ ,  $\mathcal{R}_i[t_{rs}^i, t_{rs}^i + T]$ , contains all possible trajectories of the agent from  $t_{rs}^i$  to  $t_{rs}^i + T$ . If those reachable sets are useful, then they contain all possible trajectories of two agents from  $t_c^i$  to sometime  $T_{\text{safe}} = \min(t_{rs}^j + T - \delta_i - \delta_j, t_{rs}^i + T)$  in the future based on the agent  $A_i$  clock. Therefore, the minimum distance  $d_{min}$  between two agents computed from two reachable sets is the smallest distance among all possible distances in the time




Fig. 3: Distributed Search Application Using Quadcopters.

interval  $[t_c^i, T_{safe}]$ . Consequently, the collision free guarantee is sound in the time interval  $[t_c^i, T_{safe}]$ .

## 5 Case study

The decentralized real-time safety verification for distributed CPS proposed in this paper is implemented in Java as a package called *drreach*. This package is currently integrated as a library in StarL, which is a novel platform-independent framework for programming reliable distributed robotics applications on Android [19]. StarL is specifically suitable for controlling a distributed network of robots over WiFi since it provides many useful functions and sophisticated algorithms for distributed applications. In our approach, we use the reliable communication network of StarL which is assumed to be asynchronous and peer-to-peer. There may be message dropouts and transmission delays; however, every message that an agent tries to send is eventually delivered with some time guarantees. All experimental results of our approach are reproducible and available online at: <http://www.verivital.com/rtreach/>.

We evaluate the proposed approach via a distributed search application using quadcopters<sup>3</sup> in which each quadcopter executes its search mission provided by users as a list of way-points depicted in Figure 3. These quadcopters follow the way-points to search for some specific objects. For safety reasons, they are required to work only in a specific region defined by users. In this case study, the quadcopters are controlled to operate at the same constant altitude. It has been shown from the experiments that the proposed approach is promisingly scalable as it works well for a different number of quadcopters. We choose to present in this section the experimental results for the distributed search application with eight quadcopters.

<sup>3</sup> A video recording is available at: [https://youtu.be/YC\\_7BChsIf0](https://youtu.be/YC_7BChsIf0)

The first step in our approach is locally computing the reachable set of each quadcopter using face-lifting method. The quadcopter has nonlinear motion dynamics given in Equation 2 in which  $\theta$ ,  $\phi$ , and  $\psi$  are the pitch, roll, and yaw angles,  $f = \sum_{i=1}^4 T_i$  is the sum of the propeller forces,  $m$  is the mass of the quadcopter and  $g = 9.81m/s^2$  is the gravitational acceleration constant. As the quadcopter is set to operate on a constant altitude, we have  $\ddot{z} = 0$  which yields the following constraint:  $f = \frac{mg}{\cos(\theta)\cos(\phi)}$ . Let  $v_x$  and  $v_y$  be the velocities of a quadcopter along with x- and y- axes. Using the constraint on the total force, the motion dynamics of the quadcopter can be rewritten as a 4-dimensional nonlinear ODE as depicted in Equation 3.

$$\begin{aligned} \ddot{x} &= \frac{f}{m}(\sin(\psi)\sin(\phi) + \cos(\psi)\sin(\theta)\cos(\phi)), & \dot{x} &= v_x, \\ \ddot{y} &= \frac{f}{m}(\sin(\psi)\sin(\theta)\cos(\phi) - \sin(\phi)\cos(\psi)), & \dot{v}_x &= g\tan(\theta), \\ \ddot{z} &= \frac{f}{m}\cos(\theta)\cos(\phi) - g, & \dot{y} &= v_y, \\ & & \dot{v}_y &= g\frac{\tan(\phi)}{\cos(\theta)}. \end{aligned} \quad (2) \quad (3)$$

A PID controller is designed to control the quadcopter to move from its current position to desired way-points. Details about the controller parameters can be found in the available source code. The PID controller has a control period of  $T_c = 200$  milliseconds. In every control period, the control inputs pitch ( $\theta$ ) and roll ( $\phi$ ) are computed based on the current positions of the quadcopter and the current target position (i.e., the current way-point it needs to go). Using the control inputs, the current positions and velocities given from GPS and the motion dynamics of the quadcopter, the real-time reachable set computation algorithm (Algorithm 3.1) is executed *inside* the controller. This algorithm computes the reachable set of a quadcopter from its current local time to the next  $T = 2$  seconds. The allowable run-time for this algorithm is  $T_{runtime} = 10$  milliseconds. The local safety property is verified by the real-time reachable set computation algorithm at run-time. The computed reachable set is then encoded and sent to another quadcopter. When a reachable set message arrives, the quadcopter decodes the message to reconstruct the current reachable set of the sender. The GPS error is assumed to be 2%. The time-synchronization error between the quadcopters is  $\delta = 3$  milliseconds. We want to verify in real-time: 1) local safety property for each quadcopter; 2) collision occurrence. The local safety property is defined by  $v_x \leq 500$ , i.e., the maximum allowable velocities along the x-axis of two arbitrary quadcopters are not larger than  $500m/s$ . The collision is checked using the minimum allowable distance between two arbitrary quadcopters  $d_{min} = 100$ .

Figure 4 presents a sample of a sequence of events happening in the distributed search application. One can see that each quadcopter can determine based on its local clocks if there is no collision to some known time in the future. In addition, the local safety property can also be verified at run-time. For example, in the figure, the quadcopter 1 receives a reachable set message from the quadcopter 0 which is valid from 17 : 29 : 49.075 to 17 : 29 : 51.074 of the quadcopter 0's clock. After decoding this message, taking into account the

quadcopter7 finishes computing reach set and stores the reach set  
 quadcopter6 finishes computing reach set and stores the reach set  
 quadcopter5 finishes computing reach set and stores the reach set  
 quadcopter1 finishes computing reach set and stores the reach set  
 quadcopter2 encodes its reach set to send out in 0.027134 milliseconds  
 quadcopter5 encodes its reach set to send out in 0.012657 milliseconds  
 quadcopter5 broadcasts its reach set to others  
 quadcopter2 broadcasts its reach set to others  
 quadcopter7 encodes its reach set to send out in 0.013169 milliseconds  
 quadcopter7 broadcasts its reach set to others  
 quadcopter0 encodes its reach set to send out in 0.012709 milliseconds  
 quadcopter0 broadcasts its reach set to others  
 quadcopter0 does not violate its local safety property  
 quadcopter1 encodes its reach set to send out in 0.012707 milliseconds  
 quadcopter1 broadcasts its reach set to others  
 quadcopter6 encodes its reach set to send out in 0.011081 milliseconds  
 quadcopter6 broadcasts its reach set to others  
 quadcopter1 does not violate its local safety property  
 quadcopter2 does not violate its local safety property  
 quadcopter5 does not violate its local safety property  
 quadcopter7 does not violate its local safety property  
 quadcopter6 may violates its local safety specification at time 2019-02-17 17:29:51.344  
 quadcopter7 finishes computing reach set and stores the reach set  
 quadcopter5 finishes computing reach set and stores the reach set  
 quadcopter6 finishes computing reach set and stores the reach set  
 Reach set (hull) of quadcopter0 that is valid from 2019-02-17 17:29:49.075 to 2019-02-17 17:29:51.074 of its local time is:  
 dim = 0 -> [-263.98, 1034.28]  
 dim = 1 -> [-329.46, -287.49]  
 dim = 2 -> [129.36, 301.87]  
 dim = 3 -> [30.00, 58.48]  
 Current reach set (hull) of quadcopter1 that is valid from 2019-02-17 17:29:49.386 to 2019-02-17 17:29:51.383 of its local time is:  
 dim = 0 -> [1959.99, 2040.00]  
 dim = 1 -> [-0.00, -0.00]  
 dim = 2 -> [-285.97, 395.76]  
 dim = 3 -> [-177.55, -149.38]  
 Current local time of quadcopter1 is 2019-02-17 17:29:49.423  
 Useful time for checking collision and global safety property for quadcopter1 is 1645 milliseconds  
 The received reachable set from quadcopter0 is useful  
 quadcopter1 will not collide with quadcopter0 in the next 1.645 seconds

Fig. 4: A sample of events.

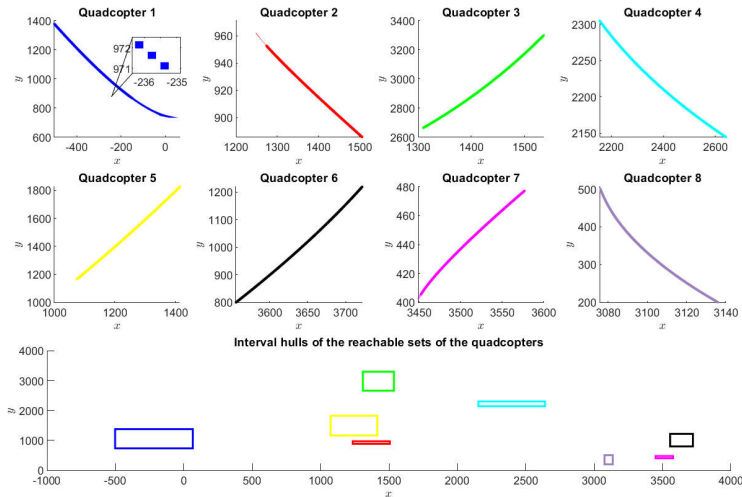


Fig. 5: One sample of the reachable sets of eight quadcopters in  $[0, 2s]$  time interval and their interval hulls.

time-synchronization error  $\delta$ , quadcopter 1 realizes that the received reachable set message is useful for checking collision for the next 1.645 seconds of its clock. After checking collision, quadcopter 1 knows that it will not collide with the quadcopter 0 in the next 1.645 seconds (based on its clock).

It should be noted that we can intuitively verify the collision occurrences by observing the intermediate reachable sets of all quadcopters and their interval hulls. The *intermediate* reachable sets of the quadcopters in every  $[0, 2s]$  time interval computed by the real-time reachable set computation algorithm (i.e., Algorithm 3.1) is described in Figure 5. The zoom plot within the figure presents a very short-time interval reachable set of the quadcopters. We note that the intermediate reachable set of a quadcopter is represented as a list of hyper-rectangles and is used for verifying the local safety property at run-time. The reachable set that is sent to another quadcopter is the interval hull of these hyper-rectangles. The intermediate reachable set cannot be transferred via a network since it is very large (i.e., hundreds of hyper-rectangles). The interval hull of all hyper-rectangles contained in the intermediate reachable set covers all possible trajectories of a quadcopter in the time interval of  $[0, 2s]$ . Therefore, it can be used for safety verification. One may question why we use the interval hull instead of using the convex hull of the reachable set since the former one results in a more conservative result. The reason is that we want to perform the safety verification online, convex hull of hundreds of hyper-rectangles is a time-consuming operation. Therefore, in the real-time setting, interval hull operation is a suitable solution. From the figure, we can see that the interval hulls of the reachable set of all quadcopters do not intersect with each other. Therefore, there is no collision occurrence (in the next 2 seconds of global time).

Time	Quad. 1	Quad. 2	Quad. 3	Quad. 4	Quad. 5	Quad. 6	Quad. 7	Quad. 8
<i>Encoding Time <math>\tau_e</math> (ms)</i>	0.058	0.055	0.0553	0.0525	0.0557	0.0583	0.0584	0.0597
<i>Decoding Time <math>\tau_d</math> (ms)</i>	0.0169	0.0193	0.0197	0.019	0.0210	0.0181	0.0177	0.022
<i>Transferring Time <math>\tau_{tf}</math> (ms)</i>	2.64	2.48	1.42	1.11	1.12	1.08	1.05	1.13
<i>Collision Checking Time <math>\tau_c</math> (ms)</i>	0.04	0.05	0.07	0.05	0.03	0.07	0.07	0.14
<i>Total Verification Time <math>VT</math> (ms)</i>	28.9363	27.9	20.6232	18.3055	18.2527	18.235	18.0223	19.1037

Table 1: The average encoding time  $\tau_e$ , decoding time  $\tau_d$ , transferring time  $\tau_{tf}$ , collision checking time  $\tau_c$  and total verification time  $VT$  of the quadcopters.

Since we implement the decentralized real-time safety verification algorithm inside the quadcopter’s controller, it is important to analyze whether or not the verification procedure affects the control performance of the controller. To reason about this, we measure the average encoding, decoding, transferring and collision checking times for all quadcopters using 100 samples which are presented in Table 1. We note that the transferring time  $\tau_{tf}$  is the average time for one message transferred from other quadcopters to the  $i^{th}$  quadcopter. It can be seen that the encoding, decoding and collision checking times at each quadcopter constitute a tiny amount of time. The total verification time is the sum of the reachable set computation, encoding, transferring, decoding and collision checking times. Note that the allowable runtime for reachable set computation

algorithm is specified by users as  $T_{runtime} = 10$  milliseconds. Therefore, the (average) total time for the safety verification procedure on each quadcopter is  $VT_i = T_{runtime} + \tau_e^i + (N - 1) \times (\tau_{tf}^i + \tau_d^i + \tau_c^i)$ , where  $i = 1, 2, \dots, N$ , and  $N$  is the number of quadcopters. As shown in the Table, the (average) total verification time for each quadcopter is small ( $< 30$  milliseconds), compared to the control period  $T_c = 200$  milliseconds. Besides, from the experiment, we observe that the computation time for the control signal of the PID controller  $\tau_{control}^i$  (not presented in the table) is also small, i.e., from 5 to 10 milliseconds. Since  $VT_i + \tau_{control}^i < T_c/4 = 50$  milliseconds, we can conclude that the verification procedure does not affect the control performance of the controller.

Interestingly, from the verification time formula, we can estimate the range of the number of agents that the decentralized real-time verification procedure can deal with. The idea is that, in each control period  $T_c$ , after computing the control signal, the remaining time bandwidth  $T_c - \tau_{control}$  can be used for verification. Let  $\bar{\tau}_e(\underline{\tau}_e)$ ,  $\bar{\tau}_{tf}(\underline{\tau}_{tf})$ ,  $\bar{\tau}_d(\underline{\tau}_d)$ ,  $\bar{\tau}_c(\underline{\tau}_c)$  be the maximum (minimum) encoding, transferring, decoding and collision checking times on a quadcopter,  $\bar{\tau}_{control}(\underline{\tau}_{control})$  be the maximum (minimum) control signal computation time for each control period  $T_c$ , then the number of agents that the decentralized real-time safety verification procedure can deal with (with assumption that the communication network works well) satisfies the following constraint:

$$\frac{T_c - \bar{\tau}_{control} - T_{runtime} - \bar{\tau}_e}{\bar{\tau}_{tf} + \bar{\tau}_d + \bar{\tau}_c} + 1 \leq N \leq \frac{T_c - \underline{\tau}_{control} - T_{runtime} - \underline{\tau}_e}{\underline{\tau}_{tf} + \underline{\tau}_d + \underline{\tau}_c} + 1. \quad (4)$$

Let consider our case study, from the Table, we assume that  $\bar{\tau}_e = 0.0597$ ,  $\underline{\tau}_e = 0.0525$ ,  $\bar{\tau}_{tf} = 2.64$ ,  $\underline{\tau}_{tf} = 1.05$ ,  $\bar{\tau}_d = 0.022$ ,  $\underline{\tau}_d = 0.0169$ ,  $\bar{\tau}_c = 0.14$ ,  $\underline{\tau}_c = 0.03$  milliseconds. Also, we assume that  $\bar{\tau}_{control} = 10$  and  $\underline{\tau}_{control} = 5$  milliseconds. We can estimate theoretically the number of quadcopters that our verification approach can deal with is  $64 \leq N \leq 168$ .

## 6 Related Work

Our work is inspired by the static and dynamic analysis of timed distributed traces [8] and the real-time reachability analysis for verified simplex design [3]. The former one proposes a sound method of constructing a global reachable set for a distributed CPS based on the recorded traces and time synchronization errors of participating agents. Then the global reachable set is used to verify a global property using Z3 [7]. This method can be considered to be a *centralized analysis* where the reachable set of the whole system is constructed and verified by one analyzer. Such a verification approach is offline which is fundamentally different from our approach as we deal with online verification in a decentralized manner. Our real-time verification method borrows the face-lifting technique developed in [3] and applies it to a distributed CPS.

Another interesting aspect of real-time monitoring for linear systems was recently published in [5]. In this work, the authors proposed an approach that combines offline and online computation to decide if a given plant model has

entered an uncontrollable state which is a state that no control strategy can be applied to prevent the plant go to the unsafe region. This method is useful for a single real-time CPS, but not a distributed CPS with multiple agents.

Additionally, there has been other significant works for verifying distributed CPS. Authors of [9, 23, 24] presented a real-time software for distributed CPS but did not perform a safety verification of individual components and a whole system. The works presented in [2, 14, 16] can be used to verify distributed CPS, but they do not consider a real-time aspect. An interesting work proposed in [21] can formally model and verify a distributed car control system against several safety objectives such as collision avoidance for an arbitrary number of cars. However, it does not address the verification problem of distributed CPS in a real-time manner. The novelty of our approach is that it can over-approximate of the reachable set of each agent whose dynamics are non-linear with a high precision degree in real-time.

The most related work to our scheme was recently introduced in [20]. The authors proposed an online verification using reachability analysis that can guarantee safe motion of mobile robots with respect to walking pedestrians modeled as hybrid systems. This work utilizes CORA toolbox [1] to perform reachability analysis while our work uses a face-lifting technique. However, this work does not consider the time-elapse for encoding, transferring and decoding the reachable set messages between each agent, which play an important role in distributed systems.

## 7 Conclusion and Future Work

We have proposed a decentralized real-time safety verification method for distributed cyber-physical systems. By utilizing the timing information and the reachable set information from exchanged reachable set messages, a sound guarantee about the safety of the whole system is obtained for each participant based on its local time. Our method has been successfully applied for a distributed search application using quadcopters built upon StarL framework. The main benefit of our approach is that it allows participants to take advantages of formal guarantees available locally in real-time to perform intelligent actions in dangerous situations. This work is a fundamental step in dealing with real-time safe motion/path planing for distributed robots. For future work, we seek to deploy this method on a real-platform and extend it to distributed CPS with heterogeneous agents where the agents can have different motion dynamics and thus they have different control periods. In addition, the scalability of the proposed method can be improved by exploiting the benefit of parallel processing, i.e., each agent handles multiple reachable set messages and checks for collision in parallel.

## Acknowledgments

The material presented in this paper is based upon work supported by the Air Force Office of Scientific Research (AFOSR) through contract number FA9550-

18-1-0122 and the Defense Advanced Research Projects Agency (DARPA) through contract number FA8750-18-C-0089. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFOSR or DARPA.

## References

1. Althoff, M.: An introduction to cora 2015. In: Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems (2015)
2. Bae, K., Krisiloff, J., Meseguer, J., Ólveczky, P.C.: Designing and verifying distributed cyber-physical systems using multirate pals: An airplane turning control system case study. *Science of Computer Programming* 103, 13–50 (2015)
3. Bak, S., Johnson, T.T., Caccamo, M., Sha, L.: Real-time reachability for verified simplex design. In: Real-Time Systems Symposium (RTSS), 2014 IEEE. pp. 138–148. IEEE (2014)
4. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Flow\*: An analyzer for non-linear hybrid systems. In: International Conference on Computer Aided Verification. pp. 258–263. Springer (2013)
5. Chen, X., Sankaranarayanan, S.: Model predictive real-time monitoring of linear systems. In: Real-Time Systems Symposium (RTSS), 2017 IEEE. pp. 297–306. IEEE (2017)
6. Dang, T., Maler, O.: Reachability analysis via face lifting. In: Hybrid Systems: Computation and Control (HSCC '98). pp. 96–109. Springer (1998), INCS 1386
7. De Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: International conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. Springer (2008)
8. Duggirala, P.S., Johnson, T.T., Zimmerman, A., Mitra, S.: Static and dynamic analysis of timed distributed traces. In: Real-Time Systems Symposium (RTSS), 2012 IEEE 33rd. pp. 173–182. IEEE (2012)
9. Eidson, J.C., Lee, E.A., Matic, S., Seshia, S.A., Zou, J.: Distributed real-time software for cyber-physical systems. *Proceedings of the IEEE* 100(1), 45–59 (2012)
10. Frehse, G., Le Guernic, C., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: Spaceex: Scalable verification of hybrid systems. In: Computer Aided Verification. pp. 379–395. Springer (2011)
11. Girard, A., Le Guernic, C., Maler, O.: Efficient computation of reachable sets of linear time-invariant systems with inputs. In: Hybrid Systems: Computation and Control, pp. 257–271. Springer (2006)
12. Henzinger, T.A.: The theory of hybrid automata. In: IEEE Symposium on Logic in Computer Science (LICS). p. 278. IEEE Computer Society, Washington, DC, USA (1996)
13. Henzinger, T.A., Ho, P.H., Wong-Toi, H.: Hytech: A model checker for hybrid systems. In: Computer aided verification. pp. 460–463. Springer (1997)
14. Johnson, T.T., Mitra, S.: Parametrized verification of distributed cyber-physical systems: An aircraft landing protocol case study. In: Cyber-Physical Systems (IC-CPS), 2012 IEEE/ACM Third International Conference on. pp. 161–170. IEEE (2012)



15. Kong, S., Gao, S., Chen, W., Clarke, E.: dreach:  $\delta$ -reachability analysis for hybrid systems pp. 200–205 (2015)
16. Kumar, P., Goswami, D., Chakraborty, S., Annaswamy, A., Lampka, K., Thiele, L.: A hybrid approach to cyber-physical systems verification. In: Proceedings of the 49th Annual Design Automation Conference. pp. 688–696. ACM (2012)
17. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21(7), 558–565 (1978)
18. Le Guernic, C., Girard, A.: Reachability analysis of hybrid systems using support functions. In: *Computer Aided Verification*. pp. 540–554. Springer (2009)
19. Lin, Y., Mitra, S.: Starl: Towards a unified framework for programming, simulating and verifying distributed robotic systems. *CoRR* abs/1502.06286 (2015), <http://arxiv.org/abs/1502.06286>
20. Liu, S.B., Roehm, H., Heinzemann, C., Lütkebohle, I., Oehlerking, J., Althoff, M.: Provably safe motion of mobile robots in human environments. In: *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*. pp. 1351–1357. IEEE (2017)
21. Loos, S.M., Platzer, A., Nistor, L.: Adaptive cruise control: Hybrid, distributed, and now formally verified. In: *International Symposium on Formal Methods*. pp. 42–56. Springer (2011)
22. Lynch, N., Segala, R., Vaandrager, F., Weinberg, H.B.: *Hybrid i/o automata*. Springer (1996)
23. Tang, Q., Gupta, S.K., Varsamopoulos, G.: A unified methodology for scheduling in distributed cyber-physical systems. *ACM Transactions on Embedded Computing Systems (TECS)* 11(S2), 57 (2012)
24. Zhang, Y., Gill, C., Lu, C.: Reconfigurable real-time middleware for distributed cyber-physical systems with aperiodic events. In: *Distributed Computing Systems, 2008. ICDCS'08. The 28th International Conference on*. pp. 581–588. IEEE (2008)