

C2E2: A Verification Tool for Stateflow Models

Parasara Sridhar Duggirala¹, Sayan Mitra², Mahesh Viswanathan¹,
and Matthew Potok²

¹ Department of Computer Science, University of Illinois at Urbana Champaign
{duggira3,vmahesh}@illinois.edu

² Department of Electrical and Computer Engineering,
University of Illinois at Urbana Champaign
mitras@illinois.edu

Abstract. Mathworks’ Stateflow is a predominant environment for modeling embedded and cyber-physical systems where control software interacts with physical processes. We present Compare-Execute-Check-Engine (C2E2)—a verification tool for continuous and hybrid Stateflow models. It checks bounded time invariant properties of models with nonlinear dynamics, and discrete transitions with guards and resets. C2E2 transforms the model, generates simulations using a validated numerical solver, and then computes reachtube over-approximations with increasing precision. For this last step it uses annotations that have to be added to the model. These annotations are extensions of proof certificates studied in Control Theory and can be automatically obtained for linear dynamics. The C2E2 algorithm is sound and it is guaranteed to terminate if the system is robustly safe (or unsafe) with respect to perturbations of guards and invariants of the model. We present the architecture of C2E2, its workflow, and examples illustrating its potential role in model-based design, verification, and validation.

1 Introduction

Cyber-physical systems (CPS) are systems that involve the close interaction between a software controller and a physical plant. The state of the physical plant evolves continuously with time and is often modeled using ordinary differential equations (ODE). The software controller, on the other hand, evolves through discrete steps and these steps influence the evolution of the physical process. This results in a “hybrid” behavior of discrete and continuous steps that makes the formal analysis of these models particularly challenging, so much so, that even models that are mathematically extremely simple are computationally intractable. In addition, many physical plants have complicated continuous dynamics that are described by nonlinear differential equations. Such plants, even without any interaction with a controlling software, are often unamenable to automated analysis.

On the other hand, the widespread deployment of CPS in safety critical scenarios like automotives, avionics, and medical devices, have made formal, automated analysis of such systems necessary. This is evident from the extensive activity in the research community [20,19,7]. Given the challenges of formally verifying CPS, the sole analysis technique that is commonly used to analyze nonlinear systems is numerical simulation. However, given the large, uncountable space of behaviors, using numerical simulations

to discover design flaws is like searching for a needle in the proverbial haystack. In this paper we present a tool C2E2 (Compare-Execute-Check-Engine) that leverages the power of numerical simulations to formally prove or disprove the safety of CPS over a bounded time interval.

Systems analyzed by C2E2 are described using Mathworks™ Stateflow™ diagrams. Stateflow is the predominant, even de facto standard, environment for designing and developing embedded and CPS both in the industry and in academia. C2E2 interprets Stateflow designs as hybrid automata [14], which is a popular mathematical model, with precise semantics, for describing CPS. The models given as input to C2E2 must be *annotated*. The annotations here are similar in spirit to code assertions and contracts used in the software verification domain. Each mode in the Stateflow diagram has to be annotated with what we call a *discrepancy function* by the user. Discrepancy functions are a generalization of several proof certificates used in control theory for analyzing convergence and divergence of trajectories. In [8] we define discrepancy functions and discuss how they can be computed automatically for a reasonably expressive class of models. C2E2 transforms the input model and compiles it with a numerical simulation library to produce a validated simulator for the model. This simulator is then used for computing increasingly precise reach set over-approximations until it proves or disproves the bounded time safety property.

Our simulation based verification approach underlying C2E2, was first presented in [8] and was subsequently used for a significant case study in [9]. The current paper outlines several enhancements to the C2E2 tool that have been made since then. First the verification algorithm presented in [8] only worked for *switched systems*¹ where the time of mode switches is explicitly given in the input. In this paper, we extend the algorithm to analyze “full hybrid automata”, i.e., continuous variables can be reset on mode switches, and mode switches take place based on enabling guards rather than explicitly given times. These theoretical improvements enable us to use C2E2 on a new set of examples. We report our experience in using C2E2 on these examples and its performance as experimental results. Next, C2E2 has been engineered to be more robust, moving from an in-house prototype, to something that can be used by the wider academic community. Finally, the tool now has a few additional features that make for a better user experience. First, it has been integrated with Stateflow, to make it useful for a wider community. Second, it can be used through a graphical user interface. And lastly, visualization tools have been added to enable users to plot various aspects of the reachable state space.

1.1 Related Work

Tools for verifying CPS vary widely based on the complexity of the complexity of the continuous dynamics. Uppaal [15], HyTech [11], and SpaceEx [10] are tools verifying timed automata, rectangular hybrid automata and linear hybrid automata respectively. Current tools available for verifying nonlinear dynamics are d/dt [2], Flow* [5] and Ariadne [3]. Typically these tools use symbolic methods for computing reachable set

¹ Switched system here refers to a design where the state of the physical plant does not change when a discrete transition is taken.

of states (or their overapproximations) from a given initial set for inferring safety of the system. Such tools provide formal soundness guarantees, however, do not provide relative completeness. Further, these tools do not analyze Stateflow models and hence a user has to specify model in the a specific input language, which requires additional learning curve for using these tools.

Given the popularity of Simulink-Stateflow framework to model CPS, there are several MATLAB based tools which verify such models. Breach [7] uses sensitivity analysis for analyzing MTL properties of systems using simulations. This analysis is sound and relatively complete for linear systems, but does not provide formal guarantees for nonlinear systems. S-Taliro [19] is a falsification engine that searches for counterexamples using Monte-Carlo techniques and hence provides only probabilistic guarantees. STRONG [6] uses robustness analysis for coverage of all executions from the initial set, using Lyapunov functions, however cannot handle nonlinear systems. C2E2, although requires additional annotations, can handle nonlinear systems specified in Stateflow and provide rigorous soundness and completeness guarantees. The simulation based verification algorithm in [8] has been extended for more general properties in [9] and to networked input output systems in [13].

2 Hybrid Models and Safety Verification

Hybrid automata is a convenient and widely used mathematical framework for modeling CPS. One of its key features is that it combines two distinct modeling styles, namely, differential equations and automata. For CPS, this enables the physical environment to be modeled by differential equations and software and communication to modeled by discrete state transitions.

Figure 1(a) shows an example of a simplified hybrid model of cardiac cell with a pacemaker created using Mathworks' Stateflow. The pacemaker has two modes or *locations* of operation: in the *stimOn* mode the cell is electrically stimulated by the pacemaker, and in *stimOff* the stimulus is absent. The continuous variables u and v model certain electrical properties of the cell. The stimulus is applied to the cell at regular time intervals as measured by the clock t . The resulting evolution of one of the continuous variables over time is shown in Figure 1(b).

Although there is no published formal semantics of Stateflow models, it is standard to consider them as hybrid automata [17,12,7]. Let us denote the set of all the variables (both continuous and discrete) in the model as the set V . This set includes a special variable loc to denote the current location. In this case, loc can be either *stimOff* or *stimOn*. The rest of the variables in V are continuous and real-valued. The set of all possible valuations of all variables in V is denoted as $val(V)$ —this defines the set of states of the model. The continuous evolution of the variables is modeled by *trajectories*. A single trajectory τ is a function $\tau : [0, t] \rightarrow val(V)$, where $t \geq 0$ is the duration τ . The state of the system at a given time t in τ is $\tau(t)$, and the value of a particular variable $v \in V$ at that state is denoted by $\tau(t).v$. A set of trajectories is specified by differential equations involving the continuous variables (see, for example, Figure 1). A trajectory τ satisfies an ordinary differential equation (ODE) $\dot{v} = f(v)$ if at each time




Fig. 1. (a) Simplified StateflowTM model of a cardiac cell and a pacemaker. (b) A simulation of the model from an initial state.

t in the domain of the trajectory, $\frac{d(\tau(t).v)}{dt} = f(\tau(t).v)$. When f is a nice² function, the ODE has a unique solution for a given initial state and a duration. With different initial states and time bounds, an ODE defines a set of trajectories.

The discrete transitions between the two locations are specified by a set A of actions (see Figure 1). An action $a \in A$ is enabled at a state whenever the state satisfies a special predicate $Guard_a$ that is associated with the action. The discrete transition from the location `stimOn` to `stimOff` is enabled only when $t \geq 5$, that is the clock has counted 5 units in `stimOn`. When the system takes a discrete transition, the new state of the system after the transition is defined by a function $Reset_a$ that maps the old state to a new state (and possibly a new location). For example, the reset function for `stimOn` to `stimOff` sets $t = 0$, $loc = stimOff$, and leaves the other continuous variables u and v unchanged. All of these components together defines the behavior of the hybrid automaton in terms of a sequence of alternating trajectories and transitions.

Definition 1. A Hybrid Automaton (HA) \mathcal{A} is a tuple $\langle V, Loc, A, \mathcal{D}, \mathcal{T} \rangle$ where

- $V = X \cup \{loc\}$ is a set of variables. Here loc is a discrete variable of finite type Loc . Valuations of loc are called locations. Each $x \in X$ is a continuous variable of type \mathbb{R} . Elements of $val(V)$ are called states.
- A is a finite set of actions or transition labels.
- $\mathcal{D} \subseteq val(V) \times A \times val(V)$ is the set of discrete transitions. A discrete transition $(\mathbf{v}, a, \mathbf{v}') \in \mathcal{D}$ is written as $\mathbf{v} \xrightarrow{a} \mathbf{v}'$. The discrete transitions are specified by finitely many guards and reset maps involving V .
- \mathcal{T} is a set of trajectories for X which is closed under suffix, prefix and concatenation (see [14] for details). For each $l \in Loc$, a set of trajectories \mathcal{T}_l for location l are specified by differential equations E_l and an invariant $I_l \subseteq val(X)$. Over any trajectory $\tau \in \mathcal{T}_l$, loc remains constant and the variables in X evolve according to E_l such that for all at each time in the domain of τ , $\tau(t)$ satisfies the invariant I_l .

An execution of a hybrid automaton \mathcal{A} records all the information (about variables) over a particular run. Formally, an *execution* is an alternating sequence of trajectories

² For example, Lipschitz continuous or smooth. A continuous function $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ is *smooth* if all its higher derivatives and partial derivatives exist and are also continuous. It has a Lipschitz constant $K \geq 0$ if for every $x_1, x_2 \in \mathbb{R}^n$, $\|f(x_1) - f(x_2)\| \leq K\|x_1 - x_2\|$.

and actions $\sigma = \tau_0 a_1 \tau_1 \dots$ where each τ_i is a closed trajectory and $\tau_i(t) \xrightarrow{a_{i+1}} \tau_{i+1}(0)$, where t is the last time point in τ_i . The *duration* of an execution is the total duration of all the trajectories. The set of all executions is denoted as $execs(\mathcal{A})$. In this paper, we only consider executions with bounded number of switches and with bounded duration. Given a set of initial states $\Theta \subseteq val(V)$, the set of *executions from* Θ are those executions in $execs(\mathcal{A})$ with their first state, $\tau_0(0)$, in Θ . The set of executions starting from Θ of duration at most T and with at most N transitions will be denoted as $execs(\mathcal{A}, \Theta, T, N)$.

Definition 2 (Safe and Unsafe). *Given a hybrid automaton \mathcal{A} with an initial set Θ , unsafe set U , time bound T , and transition bound N , it is said to be unsafe, if there exists an execution $\tau_0 a_1 \dots \tau_k \in execs(\mathcal{A}, \Theta, T, N)$ such that $\tau_k(t) \in U$, where t is the last time point in τ_k . Otherwise, \mathcal{A} is said to be safe.*

Our algorithm for safety verification as well as its analysis relies heavily on the notion of distance between continuous states and trajectories of the automaton. To state our results formally, we need to introduce a few notations first. For a vector $x \in \mathbb{R}^n$, $\|x\|$ denotes the ℓ^2 norm. For $x_1, x_2 \in \mathbb{R}^n$, $\|x_1 - x_2\|$ is the Euclidean distance between the points. For $\delta > 0$, $B_\delta(x_1) \subseteq \mathbb{R}^n$ denotes the set of points that is a closed ball of radius δ centered at x_1 . For a set $S \subseteq \mathbb{R}^n$, $B_\delta(S) = \cup_{x \in S} B_\delta(x)$. $B_\delta(S)$ expands S by δ . We will find it convenient to also define the notion of shrinking S by δ : For $\delta < 0$, and $S \subseteq \mathbb{R}^n$, $B_\delta(S) = \{x \in S \mid B_{-\delta}(x) \subseteq S\}$. For a bounded set S , a δ -cover of S is a finite collection of balls $\mathcal{X} = \{B_\delta(x_i)\}_{i=1}^m$ such that $S \subseteq \bigcup_{i=1}^m B_\delta(x_i)$. Its diameter $dia(S) \triangleq \sup_{x_1, x_2 \in S} \|x_1 - x_2\|$.

Definition 3 (Perturbing a Hybrid Automaton). *Given a hybrid automaton $\mathcal{A} = \langle V, Loc, A, \mathcal{D}, \mathcal{T} \rangle$, we define an ϵ -perturbation of \mathcal{A} as a new automaton \mathcal{A}_ϵ that has identical components as \mathcal{A} , except, (a) for each location $\ell \in Loc$, $I_\ell^{\mathcal{A}_\epsilon} = B_\epsilon(I_\ell^{\mathcal{A}})$ and (b) for each action $a \in A$, $Guard_a^{\mathcal{A}_\epsilon} = B_\epsilon(Guard_a^{\mathcal{A}})$.*

Here $I_{loc}^{\mathcal{A}}$ is the invariant of the location of a hybrid automaton \mathcal{A} and $Guard_a^{\mathcal{A}}$ denotes the guard set for action a . The definition permits $\epsilon < 0$ for perturbation of a hybrid automaton. Informally, a positive perturbation of a hybrid automaton \mathcal{A} bloats the invariants and guard sets and therefore enlarges the set of executions. A negative perturbation on the other hand, shrinks the invariants and the guards and therefore reduces the set of executions.

Definition 4 (Robust Safety and Unsafety). *Given a hybrid automaton \mathcal{A} with an initial set Θ , unsafe set U , time bound T , and bound on discrete transitions N , it is said to be robustly safe if and only if $\exists \epsilon > 0$, such that \mathcal{A}_ϵ , with initial set Θ_ϵ , unsafe set U_ϵ , time bound T , and transition bound N is safe. It is said to be robustly unsafe if and only if $\exists \epsilon < 0$ such that \mathcal{A}_ϵ , with initial set Θ_ϵ , unsafe set U_ϵ , time bound T , and transition bound N , is unsafe.*

For safety verification C2E2 expects the users to provide annotations for the ODEs defining the trajectories of the hybrid automaton in question. These annotations are called discrepancy function. For a differential equation $\dot{x} = f(x, t)$, the general definition of discrepancy function is given in [8]. In this paper, we consider a special form of discrepancy function given in Definition 5.

Definition 5. Given a differential equation $\dot{x} = f(x)$, the tuple $\langle K, \gamma \rangle$ is called an exponential discrepancy function for the dynamics, if and only if for any two trajectories τ_1, τ_2 satisfying the differential equation, it holds that:

$$\|\tau_1(t) - \tau_2(t)\| \leq K \|\tau_1(0) - \tau_2(0)\| e^{\gamma t} \quad (1)$$

We call K as the multiplicity factor and γ as the exponential factor of the annotation.

Along with the input models, we expect the user to specify the discrepancy function by providing values for $\langle K, \gamma \rangle$. Also, if the differential equation is Lipschitz continuous with constant L , then $\langle K, \gamma \rangle = \langle 1, L \rangle$ would be a valid annotation. In Figure 1, the annotations for both the locations are provided as $\langle K, \gamma \rangle = \langle 3.8, -0.2 \rangle$.

3 Verifying Hybrid Systems from Simulations

We give an overview of the verification algorithm implemented in C2E2. In brief, the algorithm generates a cover of the initial set (a collection of regions) and then performs four steps repeatedly until it reaches a safe/unsafe decision for each region in the cover or its refinement. In making this decision, it generates a simulation from the center of the region and then bloats this simulation by a factor computed from the given annotation to compute an over-approximation of the reachable states. If the over-approximation decides safe/unsafe then it moves on to another region in the cover, otherwise, it refines the cover to obtain better over-approximation of the reachable states. These operations are performed by several subroutines which we describe next.

3.1 Building Blocks

The first building block for the algorithm is a subroutine called *valSim* that generates validated simulations for the individual dynamical systems or locations of the hybrid automaton. Given a trajectory τ starting from a given state, *valSim* subroutine computes overapproximation of τ with specific error.

Definition 6 (Validated Simulation). Given an error bound $\epsilon > 0$, a time step $h > 0$, a time bound $T = (k + 1)h$, and an initial state x_0 , an (x_0, ϵ, h, T) -simulation of the differential equation $\dot{x} = f(x)$ is a sequence of sets of continuous states $\rho = R_0, R_1, \dots, R_k$ such that (a) for any i in the sequence $\text{dia}(R_i) \leq \epsilon$, and (b) for any time t in the interval $[ih, (i+1)h]$, the solution from x_0 at time t is in R_i , i.e., $\tau(t) \in R_i$.

The subroutine *valSim*(x_0, h, T, f) returns a tuple $\langle \rho, \epsilon \rangle$ such that ρ is an (x_0, ϵ, h, T) -simulation. In C2E2, validated simulation engines such as VNODE-LP [18] and CAPD [4] are used for implementing *valSim*. For the completeness of our algorithm, we require that the error ϵ can be made arbitrarily small by decreasing h . In systems with finite precision arithmetic, simulation engines produce accurate simulations up to the order of 10^{-7} , and also there are libraries supporting arbitrary precision integration for some differential equations. Next, the *computeReachTube* subroutine: it uses simulations to compute over-approximations of a set of trajectories from a set of initial states.

Definition 7 (Overapproximate Reach Tube). For a set of initial states S , error bound $\epsilon > 0$, time step $h > 0$, and $T = (k + 1)h$, an (S, ϵ, h, T) -reachtube of the differential equation $\dot{x} = f(x)$ is a sequence $\psi = R_0, R_1, \dots, R_k$ such that (a) for any i in the sequence $\text{dia}(R_i) \leq \epsilon$, (b) for any trajectory starting τ from S and for each time $t \in [ih, (i + 1)h]$, $\tau(t) \in R_i$.

If S is large and ϵ is small then the strict inclusion may preclude a (S, ϵ, h, T) -reachtube from existing. To compute reachtubes from a compact set S centered at x_0 , `computeReachTube` performs the following three steps:

1. $\langle \rho, \epsilon_1 \rangle \leftarrow \text{valSim}(x_0, h, T, f)$, let $\rho = R_0, \dots, R_k$.
2. $\epsilon_2 = \sup\{K\|x_1 - x_2\|e^{\gamma t} \mid x_1, x_2 \in S, t \in [0, T]\}$, where $\langle K, \gamma \rangle$ is the annotation for the dynamics given by f .
3. $\psi = B_{\epsilon_2}(R_1), \dots, B_{\epsilon_2}(R_k)$.

From Definition 5 it follows that ψ is a $(S, \epsilon_1 + \epsilon_2, h, T)$ -reachtube. It can be shown [8] that $\epsilon \rightarrow 0$ as $\delta \rightarrow 0$ and $h \rightarrow 0$. In summary, the subroutine call `computeReachTube` (S, h, T, f) returns $\langle \psi, \epsilon \rangle$ using the above steps. To provide formal guarantees from reachtube over-approximations, we need to distinguish when a given reachtube *must* satisfy a predicate P from when it *may* satisfy it. The next subroutine `tagRegion` tags each region in a reachtube with respect to a given predicate.

Definition 8 (Tagging). Given two sets $R, P \subseteq \mathbb{R}^n$ the subroutine `tagRegion` (R, P) returns *must*, *may* or \perp such that, (a) if $R \subseteq P$, then return *must*, (b) if $R \cap P \neq \emptyset$ and $R \not\subseteq P$ then return *may*, and (c) otherwise ($R \cap P = \emptyset$), return $\text{tag} = \perp$.

The above subroutines are used for over-approximating reach sets and deciding safety for individual differential equations of individual locations of a hybrid automaton. In order to reason about invariants, guards, and resets which are essential for capturing the hybrid behavior of mode switches, we have to (a) detect the reachable states that satisfy the respective location invariants and (b) identify the states from which location switches or transitions can occur. The next subroutine `invariantPrefix` (ψ, S) takes a reachtube and a set S and returns the longest contiguous prefix of ψ that intersects with S . This is later used in solving problem (a).

Definition 9 (Invariant Prefix). Given a reachtube $\psi = R_0, \dots, R_k$ and a set S , `invariantPrefix` (ψ, S) returns the longest sequence $\phi = \langle R_0, \text{tag}_0 \rangle \dots \langle R_m, \text{tag}_m \rangle$, such that $\forall 0 \leq i \leq m$, $\text{tag}_i = \text{must}$ when $\forall j \leq i$, $\text{tagRegion}(R_j, S) = \text{must}$ and $\text{tag}_i = \text{may}$ when $\forall j \leq i$, $\text{tagRegion}(R_j, S) \neq \perp$ and $\exists l \leq i$, $\text{tagRegion}(R_l, S) = \text{may}$. That is, if $m < k$, then $\text{tagRegion}(R_{m+1}, S) = \perp$.

Intuitively, a region R_i is tagged *must* in an invariant prefix, if all the regions before R_i (including itself) are contained within the set S . It is tagged *may* if all the regions before it have nonempty intersection with S , and at least one of them (including itself) is not contained within the set S . Given a reachtube ψ from a set Q , $\phi = \text{invariantPrefix}(\psi, \text{Inv}_{loc})$ returns the over-approximation of the valid set of trajectories from Q that respect the invariant of the location loc of the hybrid automaton. If a region R_i in ϕ is tagged *must*, then there exists at least one trajectory from Q that

can reach R_i . Also, the set of all reachable states from Q that satisfy the invariant are contained in ϕ . Subroutine *checkSafety* checks whether such an invariant prefix ϕ is safe with respect to an unsafe set U . It defined as:

$$checkSafety(\phi, U) = \begin{cases} safe, & \text{if } \forall \langle R, \cdot \rangle \in \phi, \perp = tagRegion(R, U) \\ unsafe, & \text{if } \exists \langle R, must \rangle \in \phi, must = tagRegion(R, U) \\ unknown, & \text{otherwise} \end{cases}$$

Notice that ϕ is inferred to be *safe* only when all the regions in the invariant prefix are tagged \perp with respect to U (i.e. empty intersection). It is *unsafe* only when there is a *must* region in the invariant prefix is contained within U . This is the core of the soundness argument for the verification algorithm, i.e., if *checkSafety*(ϕ, U) returns *safe* or *unsafe*, then indeed that is the correct answer for the set of initial states covered by ϕ .

The final subroutine *nextRegions* computes over-approximation of the reachable states that serve as the initial states in a new location after a transition to it. A discrete transition a is enabled at the instance trajectory τ satisfies the guard condition $Guard_a$ and the state after the discrete transition is obtained by applying the reset map $Reset_a$. Given a sequence of tagged regions, the subroutine *nextRegions* returns the tagged set of reachable regions after a discrete transition.

Definition 10 (Next Regions). Given $\phi = \langle R_0, tag_0 \rangle, \dots, \langle R_m, tag_m \rangle$, a sequence of tagged regions, the subroutine *nextRegions*(ϕ) returns a set of tagged regions \mathbf{R} . $\langle R', tag' \rangle \in \mathbf{R}$ if and only if there exists an action a of the automaton and a region R_i in ϕ such that $R' = Reset_a(R_i)$ and one of the following conditions hold:

- (a) $R_i \subseteq Guard_a, tag_i = must, tag' = must$.
- (b) $R_i \cap Guard_a \neq \emptyset, R_i \not\subseteq Guard_a, tag_i = must, tag' = may$.
- (c) $R_i \cap Guard_a \neq \emptyset, tag_i = may, tag' = may$.

A tagged region $\langle R', tag' \rangle \in \mathbf{R}$ is labeled *must* only when the region R_i is a *must* region and is contained within the $Guard_a$. In all other cases, the region is tagged *may* when $R_i \cap Guard_a \neq \emptyset$. This ensures that a regions tagged *must* are indeed reachable after the discrete transition, and all the regions tagged *may* contain the reachable states after the discrete transition.

3.2 Verification Algorithm

The C2E2 algorithm for verification is given in Algorithm 1. Lines 4 - 17 implement the main loop that computes a cover of the initial set, computes the overapproximation of reachtube, checks safety, and refines the cover if needed. The subroutine *taggedCover* (line 3), first computes a cover of S (Θ in the first iteration) with sets of δ diameter and tags them (with *tagRegion*) with respect to S . The resulting tagged sets are collected in \mathcal{X} . We add additional attributes to each tagged region: *Time* tracks the time of the trajectory leading to a region, *Loc* tracks its current location, and *Switches* tracks the number of discrete transitions taken. Although not explicitly mentioned in the algorithm, we update the tags for regions in the reachtube based on the time taken by

trajectories and the discrete transitions encountered during verification. The algorithm checks whether all the executions of hybrid automaton from the regions in the cover of the initial set are safe or unsafe. If it is safe, then the region is removed from the set S (line 13); if it is unsafe, then the algorithm returns unsafe (line 10); and if it is neither then the cover is refined (line 15).

For each of the regions in the tagged δ -cover of the initial set, the inner loop (lines 6 - 11) computes over-approximations of the reachable states. The subroutine *computeReachTube* (line 7) computes an overapproximation of all the trajectories starting from the tagged region in the initial set cover. Subroutine *invariantPrefix* (line 8) computes ϕ , an overapproximation of the valid set of trajectories which respect the invariant of the current location. Safety of ϕ is checked using the subroutine *checkSafety* (line 9). The regions in ϕ where discrete transitions are enabled and the initial states for the trajectories in the next location are computed in subroutine *nextRegions* (line 10). This loop continues until either the time bound for verification or the bound for number of discrete transitions is satisfied.

The reachtubes computed in lines 7, 8 are an overapproximation of all the reachable states of hybrid automaton \mathcal{A} , given in Lemma 1. The regions in the reachtubes are *must* only when they satisfy the invariant of each location and completely contained within the guard set it follows that these regions are reachable by at least one execution of \mathcal{A} , given as Lemma 2. For the set of regions \mathcal{R} tagged *may*, it follows that $\epsilon = \max\{dia(R) | R \in \mathcal{R}\}$ bloating of the invariants and guard sets would ensure that these regions are also reachable, given in Lemma 3. As the reachtubes can be made arbitrarily precise by decreasing the time step and the initial partitioning, it follows that given any ϵ , the precision of reachtubes can be bounded by ϵ by using small values for δ and h . Lemma 1 helps in proving soundness and Lemmas 2 and 3 help in proving relative completeness. This guarantees that if the system is robustly safe, then the algorithm terminates and returns safe, and if the system is robustly unsafe, then the algorithm terminates and returns unsafe, given in Theorem 1.

Lemma 1. *All the regions R tagged *may* or *must* in ϕ (line 7 of algorithm 1) contains the reachable set of states of hybrid automaton \mathcal{A} starting from Θ within T time and N discrete transitions.*

Lemma 2. *If a region R is tagged *must* in ϕ (line 7 of algorithm 1), then there exists at least one execution of \mathcal{A} from the initial set Θ that reaches R within T time and N discrete transitions.*

Lemma 3. *Let \mathcal{R} , be the set of all regions tagged *may* in ϕ (line 7 of algorithm 1) and $\epsilon = \max\{dia(R) | R \in \mathcal{R}\}$. Given any region $R \in \mathcal{R}$, there exists at least one execution of \mathcal{A}_ϵ from the initial set Θ_ϵ that reaches R within T time and N discrete transitions.*

Theorem 1 (Soundness and Relative Completeness). *Given initial set Θ , unsafe set U , time bound T , bound on discrete transitions N , and hybrid automaton \mathcal{A} , if the algorithm 1 returns safe or unsafe, then the system \mathcal{A} is safe or unsafe. The algorithm will always terminate whenever the system is either robustly safe or robustly unsafe.*

```

input :  $\mathcal{A}, \Theta, U, T, N$ 
output: System is safe or unsafe
1  $S \leftarrow \Theta; h \leftarrow h_0; \delta \leftarrow \delta_0;$ 
2 while  $S \neq \emptyset$  do
3    $\mathcal{X} \leftarrow \text{taggedCover}(S, \delta);$ 
4   for  $elem \in \mathcal{X}$  do
5      $Q_C \leftarrow \{elem\};$ 
6     for  $e \in Q_C \wedge e.Time < T \wedge e.Switches < N$  do
7        $\langle \psi, \epsilon \rangle \leftarrow \text{computeReachTube}(e.R, e.Loc, h, T - e.Time);$ 
8        $\phi \leftarrow \text{invariantPrefix}(\psi, Inv_{e.Loc});$ 
9        $result \leftarrow \text{checkSafety}(\phi, U);$ 
10      if  $result == safe$  then  $Q_C \leftarrow Q_C \cup \text{nextRegions}(\phi)$  else if
         $result == unsafe \wedge elem.tag == must$  then return unsafe else break
11    end
12    if  $result == safe$  then
13       $S \leftarrow S \setminus elem.R;$ 
14    else
15       $\delta \leftarrow \delta/2; h \leftarrow h/2;$ 
16    end
17  end
18 end
19 return safe;

```

Algorithm 1. Algorithm for safety verification of hybrid automata using simulations and annotations.

4 C2E2: Internals and User Experience

4.1 Architecture of C2E2

The architecture for C2E2 is shown in Figure 2. The front end parses the input models, connects to the verification engine, provides a property editor and a plotter. It is developed in Python and vastly extends the Hylink parser [17] for Stateflow models. The verification algorithm (Algorithm 1) is implemented in C++. The front end parses the input model file (.mdl or .hyxml) into an intermediate format and generates the simulation code. The properties are obtained from the input file or from the user through the front end's GUI. The simulation code is compiled using a validated simulation engine provided by Computer Assisted Proofs in Dynamic Groups (CAPD) library [4]. This compiled code and the property are read by the C2E2 verification algorithm which also uses the GLPK libraries. The verification result and the computed reachable set are read by the front end for display and visualization. This modular architecture allows us to extend the functionality of the tool to new types of models (such as Simulink, DAEs), different simulation engines (for example, Boost, VNODE-LP), and alternative checkers (such as Z3).




Fig. 2. Architecture of C2E2

4.2 Models, Properties, and Annotations

C2E2³ takes as input annotated Stateflow models (such as in Figure 1(a)) with ODEs (possibly nonlinear) and discrete transitions defined by guards and resets. The guards have to be conjunctions of polynomial predicates over the state variables and the reset maps have to be polynomial real-valued functions. The properties can be specified in the .hyxml model files or using the GUI. C2E2 can verify bounded time safety properties specified by a time bound, a polyhedral set of initial states and another set of unsafe states.

Annotations. The user has to write annotations for each ODE in the model. This is done by specifying the multiplicity factor K and the exponential factor γ as comments in the Stateflow model, shown in Figure 1(a). For a broad range of nonlinear systems such annotations can be found. We illustrate with some examples.

Example 1. Consider a linear system $\dot{x} = Ax$, where all the eigenvalues of matrix A is nonzero. Let λ_m be the maximum among the real parts of all the eigenvalues of A . Then, for any two trajectories τ_1 and τ_2 of the linear system $\|\tau_1(t) - \tau_2(t)\| \leq \|A\|e^{\lambda_m t} \|\tau_1(0) - \tau_2(0)\|$ is an annotation. Here the matrix norm $\|A\|$ is defined as $\sup\{x^T Ax \mid \|x\| = 1\}$ and it can be computed using semidefinite programming. The input format would be $K = \|A\|$ and $\gamma = \lambda_m$.

Example 2. Consider the differential equation in *stimOn* mode of the cardiac cell example: $\dot{u} = (0.1 - u)(u - 1)u - v$ and $\dot{v} = u - 2v$. By computing the maximum eigenvalues of the Jacobian matrix of the differential equation and the maximum norm of the contraction metric [16], we get that $\|\tau_1(t) - \tau_2(t)\| \leq 3.8e^{-0.2t} \|\tau_1.(0) - \tau_2.(0)\|$ as an annotation. The input is specified as $K = 3.8$ and $\gamma = -0.2$.

Example 3. Consider the differential equation $\dot{x} = 1 + x^2y - 2.5x$; $\dot{y} = 1.5x - x^2y - y$. By analyzing the auxiliary system (x_1, y_1) and (x_2, y_2) , using the incremental stability [1], we have that

$$\frac{d}{dt}[(x_1 - x_2)^2 + 2(x_1 - x_2)(y_1 - y_2) + (y_1 - y_2)^2] = -2(x_1 - x_2 + y_1 - y_2)^2 < 0.$$

³ <https://publish.illinois.edu/c2e2-tool/>

Therefore, for trajectories τ_1 and τ_2 of the differential equations it hence follows that $(\tau_1(t).x - \tau_2(t).x)^2 + 2(\tau_1(t).x - \tau_2(t).x)(\tau_1(t).y - \tau_2(t).y) + (\tau_1(t).y - \tau_2(t).y)^2 \leq (x_1 - x_2)^2 + 2(x_1 - x_2)(y_1 - y_2) + (y_1 - y_2)^2$ where (x_1, y_1) and (x_2, y_2) are the initial states of τ_1 and τ_2 respectively. The function $\|\tau_1(t) - \tau_2(t)\| \leq 2\|\tau_1(0) - \tau_2(0)\|e^{0 \times t}$ is an annotation and is specified as $K = 2, \gamma = 0$.

4.3 User Experience

In this section, we discuss the C2E2 interface for handling verification, properties and visualization. The users can add, edit, copy, delete or verify several properties. As each property is edited, the *smart parser* provides real-time feedback about syntax errors and unbounded initial sets (Figure 3(b)). Once properties are edited the verifier can be launched. Visual representation of reachable states and locations can aid debugging process. To this end, we have integrated a visualizer into C2E2 for plotting the projections of the reachable states. Once a property has been verified, the user can plot the valuations of variables against time or valuations of pairs of variables (phase plots). The unsafe set is projected on the set of plotting variables. The property parser and visualizer uses the Parma Polyhedra Library⁴ and matplotlib⁵. Example plot for the cardiac cell is shown in Figure 3(c).




Fig. 3. Figure showing a snippet of cardiac cell model in (a), property dialog for specifying properties in (b), plot of reachable set for cardiac cell model in (c)

4.4 Stateflow Model Semantics

The annotated Stateflow models when given as input are interpreted as hybrid automaton as C2E2. Nondeterminism, which is allowed in hybrid automata framework is prohibited in Stateflow models. All the discrete transitions in Stateflow models are deterministic and are interpreted as “urgent” i.e. a transition is taken by the system *as soon as* it is enabled. In our front end, we construct hybrid automaton for Stateflow models respecting the “urgent” semantics. Under such interpretation, the guard sets are only allowed

⁴ <http://bugseng.com/products/ppl/>

⁵ <http://matplotlib.org/>

to be hyperplanes. In general, the verification algorithm in Section 3 may not terminate for such guard conditions. We therefore use a heuristic and verify an ϵ perturbed model of the Stateflow model to ensure termination of verification algorithm.

4.5 Experiments

Simulation based verification approach for annotated models has been demonstrated to outperform other verification tool such as Flow* and Ariadne in [8]. In this paper, we present the verification results for some of the nonlinear and linear hybrid automata benchmarks in Table 1. The annotations for each of these benchmarks have been obtained by procedures given in Section 4.2. All the experiments have been performed on Intel i-7 Quad core processor with 8GB ram running Ubuntu 11.10.

Table 1. Experimental Results for benchmark examples. Vars: Number of Variables, Num. Loc. : Number of discrete locations in hybrid automata, TH: Time Horizon for Verification, VT (sec) : Verification time for C2E2 in seconds, Result: Verification result of C2E2.

Benchmark	Vars.	Num. Loc.	TH	VT (sec)	Result
Cardiac Cell	3	2	15	17.74	safe
Cardiac Cell	3	2	15	1.91	unsafe
Nonlinear Navigation	4	4	2.0	124.10	safe
Nonlinear Navigation	4	4	2.0	4.94	unsafe
Inverted Pendulum	2	1	10	1.27	safe
Inverted Pendulum	2	1	10	1.32	unsafe
Navigation Benchmark	4	4	2.0	94.35	safe
Navigation Benchmark	4	4	2.0	4.74	unsafe

This early termination strategy for unsafe behavior of the system (Algorithm 1, line 10) when the system is unsafe is reflected in Table 1. On standard examples C2E2 can successfully verify these systems within the order of minutes and also handle non-linear differential equations with trigonometric functions of inverted pendulum.

5 Conclusions and Future Work

C2E2 presented in this paper is a tool for verifying a broad class of hybrid and dynamical systems models. It uses validated simulations and model annotations to prove the most commonly encountered type of properties, namely bounded-time invariants. It can handle models created using the Stateflow environment that is the de facto standard in embedded control design and implementation. The improvements presented in this paper beyond the version of [8], include the complete support for hybrid models implemented in a new algorithm and the supporting theory, the new user interface for editing

properties, and the reachtube plotting function. The tool is freely available for academic and research use from <https://publish.illinois.edu/c2e2-tool/>.

Our future plans include implementation of features to support temporal precedence properties [9] and compositional reasoning [13,12]. Another avenue of work leverages the “embarrassing parallelism” in the simulation-based approach. We anticipate that the C2E2’s architecture and its open interfaces, for example, the .hyxml input format, text-based representation of reachtubes, will support research and education in embedded and hybrid systems community by helping explore new ideas in modeling, verification, synthesis, and testing.

Acknowledgements. The authors were supported by the National Science Foundation research grant CSR 1016791.

References

1. Angeli, D.: A Lyapunov approach to incremental stability properties. *IEEE Transactions on Automatic Control* (2000)
2. Asarin, E., Dang, T., Maler, O.: The d/dt tool for verification of hybrid systems. In: Brinksma, E., Larsen, K.G. (eds.) *CAV 2002*. LNCS, vol. 2404, pp. 365–370. Springer, Heidelberg (2002)
3. Balluchi, A., Casagrande, A., Collins, P., Ferrari, A., Villa, T., Sangiovanni-Vincentelli, A.: Ariadne: a framework for reachability analysis of hybrid automata. In: *International Symposium on Mathematical Theory of Networks and Systems, MNTS (2006)*
4. CAPD. Computer assisted proofs in dynamic groups, <http://capd.ii.uj.edu.pl/index.php>
5. Chen, X., Abraham, E., Sankaranarayanan, S.: Flow*: An analyzer for non-linear hybrid systems. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 258–263. Springer, Heidelberg (2013)
6. Deng, Y., Rajhans, A., Julius, A.A.: STRONG: A trajectory-based verification toolbox for hybrid systems. In: Joshi, K., Siegle, M., Stoelinga, M., D’Argenio, P.R. (eds.) *QEST 2013*. LNCS, vol. 8054, pp. 165–168. Springer, Heidelberg (2013)
7. Donzé, A.: Breach, A toolbox for verification and parameter synthesis of hybrid systems. In: Touili, T., Cook, B., Jackson, P. (eds.) *CAV 2010*. LNCS, vol. 6174, pp. 167–170. Springer, Heidelberg (2010)
8. Duggirala, P.S., Mitra, S., Viswanathan, M.: Verification of annotated models from executions. In: *International Conference on Embedded Software, EMSOFT (2013)*
9. Duggirala, P.S., Wang, L., Mitra, S., Viswanathan, M., Muñoz, C.: Temporal precedence checking for switched models and its application to a parallel landing protocol. In: Jones, C., Pihlajasaari, P., Sun, J. (eds.) *FM 2014*. LNCS, vol. 8442, pp. 215–229. Springer, Heidelberg (2014)
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: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 379–395. Springer, Heidelberg (2011)
11. Henzinger, T.A., Ho, P.-H., Wong-Toi, H.: HyTech: A model checker for hybrid systems. In: Grumberg, O. (ed.) *CAV 1997*. LNCS, vol. 1254, pp. 460–463. Springer, Heidelberg (1997)
12. Huang, Z., Fan, C., Mereacre, A., Mitra, S., Kwiatkowska, M.: Invariant verification of non-linear hybrid automata networks of cardiac cells. In: Biere, A., Bloem, R. (eds.) *CAV 2014*. LNCS, vol. 8559, pp. 373–390. Springer, Heidelberg (2014)

13. Huang, Z., Mitra, S.: Proofs from simulations and modular annotations. In: International Conference on Hybrid Systems: Computation and Control, pp. 183–192 (2014)
14. Kaynar, D.K., Lynch, N., Segala, R., Vaandrager, F.: The Theory of Timed I/O Automata. Synthesis Lectures on Computer Science. Morgan Kaufmann (November 2005)
15. Larsen, K.G., Pettersson, P., Yi, W.: Uppaal in a nutshell. International Journal on Software Tools for Technology Transfer (STTT) 1(1), 134–152 (1997)
16. Lohmiller, W., Slotine, J.J.E.: On contraction analysis for non-linear systems. Automatica (1998)
17. Manamcheri, K., Mitra, S., Bak, S., Caccamo, M.: A step towards verification and synthesis from simulink/stateflow models. In: International Conference on Hybrid Systems: Computation and Control, HSCC (2011)
18. Nediaklov, N.: VNODE-LP: Validated solutions for initial value problem for ODEs. Technical report, Department of Computing and Software, McMaster University (2006)
19. Nghiem, T., Sankaranarayanan, S., Fainekos, G., Ivancic, F., Gupta, A., Pappas, G.: Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In: International Conference on Hybrid Systems: Computation and Control HSCC (2010)
20. Zou, L., Zhan, N., Wang, S., Franzle, M., Qin, S.: Verifying simulink diagrams via a hybrid hoare logic prover. In: International Conference on Embedded Software EMSOFT (2013)