SAW: A Tool for Safety Analysis of Weakly-Hard Systems

Chao Huang¹, Kai-Chieh Chang², Chung-Wei Lin², and Qi Zhu¹

¹ Northwestern University, Evanston, USA
{chao.huang,qzhu}@northwestern.edu

² National Taiwan University, Taipei, Taiwan
551100kk@gmail.com, cwlin@csie.ntu.edu.tw

Abstract. We introduce SAW, a tool for safety analysis of weakly-hard systems, in which traditional hard timing constraints are relaxed to allow bounded deadline misses for improving design flexibility and runtime resiliency. Safety verification is a key issue for weakly-hard systems, as it ensures system safety under allowed deadline misses. Previous works are either for linear systems only, or limited to a certain type of nonlinear systems (e.g., systems that satisfy exponential stability and Lipschitz continuity of the system dynamics). In this work, we propose a new technique for infinite-time safety verification of general nonlinear weakly-hard systems. Our approach first discretizes the safe state set into grids and constructs a directed graph, where nodes represent the grids and edges represent the reachability relation. Based on graph theory and dynamic programming, our approach can effectively find the safe initial set (consisting of a set of grids), from which the system can be proven safe under given weakly-hard constraints. Experimental results demonstrate the effectiveness of our approach, when compared with the state-of-the-art. An open source implementation of our tool is available at <https://github.com/551100kk/SAW>. The virtual machine where the tool is ready to run can be found at <https://www.csie.ntu.edu.tw/~r08922054/SAW.ova>.

Keywords: Weakly-hard systems · Safety verification · Graph theory

1 Introduction

Hard timing constraints, where deadlines should always been met, have been widely used in real-time systems to ensure system safety. However, with the

This work is supported by the National Science Foundation awards 1834701, 1834324, 1839511, 1724341, and the Office of Naval Research grant N00014-19-1-2496. It is also supported by the Asian Office of Aerospace Research and Development (AOARD), jointly with the Office of Naval Research Global (ONRG), award FA2386-19-1-4037, the Taiwan Ministry of Education (MOE) grants NTU-107V0901 and NTU-108V0901, the Taiwan Ministry of Science and Technology (MOST) grants MOST-108-2636-E-002-011 and MOST-109-2636-E-002-022.

C. Huang and K.-C. Chang—Contributed equally.

© The Author(s) 2020

S. K. Lahiri and C. Wang (Eds.): CAV 2020, LNCS 12224, pp. 543–555, 2020.

https://doi.org/10.1007/978-3-030-53288-8_26

Fig. 1. A weakly-hard system with perfect sensors and actuators.

rapid increase of system functional and architectural complexity, hard deadlines have become increasingly pessimistic and often lead to infeasible designs or over provisioning of system resources [16, 20, 21, 32]. The concept of weakly-hard systems are thus proposed to relax hard timing constraints by allowing occasional deadline misses [2, 11]. This is motivated by the fact that many system functions, such as some control tasks, have certain degrees of robustness and can in fact tolerate some deadline misses, as long as those misses are bounded and dependably controlled. In recent years, considerable efforts have been made in the research of weakly-hard systems, including schedulability analysis [1, 2, 5, 12–14, 19, 25, 28, 30], opportunistic control for energy saving [18], control stability analysis and optimization [8, 10, 22, 23, 26], and control-schedule co-design under possible deadline misses [3, 6, 27]. Compared with hard deadlines, weakly-hard constraints can more accurately capture the timing requirements of those system functions that tolerate deadline misses, and significantly improve system feasibility and flexibility [16, 20]. Compared with soft deadlines, where any deadline miss is allowed, weakly-hard constraints could still provide deterministic guarantees on system safety, stability, performance, and other properties under formal analysis [17, 29].

A common type of weakly-hard model is the (m, K) constraint, which specifies that among any K consecutive task executions, at most m instances could violate their deadlines [2]. Specifically, the high-level structure of a (m, K) -constrained weakly-hard system is presented in Fig. 1. Given a sampled-data system $\dot{x} = f(x, u)$ with a sampling period $\delta > 0$, the system samples the state x at the time $t = i\delta$ for $n = 0, 1, 2, \dots$, and computes the control input u with function $\pi(x)$. If the computation completes within the given deadline, the system applies u to influence the plant's dynamics. Otherwise, the system stops the computation and applies zero control input. As aforementioned, the system should ensure the control input can be successfully computed and applied within the deadline for at least $K - m$ times over any K consecutive sampling periods.

For such weakly-hard systems, a natural and critical question is whether the system is safe by allowing deadline misses defined in a given (m, K) constraint.

There is only limited prior work in this area, while nominal systems have been adequately studied [4, 9, 15, 31]. In [8], a weakly-hard system with linear dynamic is modeled as a hybrid automaton and then the reachability of the generated hybrid automaton is verified by the tool SpaceEx [9]. In [7], the behavior of a linear weakly-hard system is transformed into a program, and program verification techniques such as abstract interpretation and SMT solvers can be applied.

In our previous work [17], the safety of nonlinear weakly-hard systems are considered for the first time. Our approach tries to derive a safe initial set for any given (m, K) constraint, that is, starting from any initial state within such set, the system will always stay within the same safe state set under the given weakly-hard constraint. Specifically, we first convert the infinite-time safety problem into a finite one by finding a set satisfying both *local safety* and *inductiveness*. The computation of such valid set heavily lies on the estimation of the system state evolution, where two key assumptions are made: 1) The system is exponentially stable under nominal cases without any deadline misses, which makes the system state contract with a constant decay rate; 2) The system dynamics are Lipschitz continuous, which helps bound the expansion under a deadline miss. Based on these two assumptions, we can abstract the safety verification problem as a one-dimensional problem and use linear programming (LP) to solve it, which we call *one-dimension abstraction* in the rest of the paper.

In practice, however, the assumptions in [17] are often hard to satisfy and the parameters of exponential stability are difficult to obtain. In addition, while the scalar abstraction provides high efficiency, the experiments demonstrate that the estimation is always over conservative. In this paper, we go one step further and present a new tool SAW for infinite-time safety verification of nonlinear weakly-hard systems **without any particular assumption on exponential stability and Lipschitz bound**, and try to be less conservative than the scalar abstraction. Formally, the problem solved by this tool is described as follows:

Problem 1. Given an (m, K) weakly-hard system with nonlinear dynamics $\dot{x} = f(x, u)$, sampling period δ , and safe set X , find a safe initial set X_0 , such that from any state $x(0) \in X_0$, the system will always be inside X .

To solve this problem, we first discretize the safe state set X into grids. We then try to find the grid set that satisfies both local safety and inductiveness. For each property, we build a directed graph, where each node corresponds to a grid and each directed edge represents the mapping between grids with respect to reachability. We will then be able to leverage graph theory to construct the initial safe set. Experimental results demonstrate that our tool is effective for general nonlinear systems.

2 Algorithms and Tool Design

The schematic diagram of our tool SAW is shown in Fig. 2. The input is a model file that specifies the system dynamics, sampling period, safe region and other parameters, and a configuration file of Flow* [4] (which is set by default but can

Fig. 2. The schematic diagram of SAW.

Algorithm 1: Overall algorithm of SAW

Data: Dynamic system f with safe state region X , the control law π , weakly-hard constraint (m, K) , sampling period δ

Result: Safe initial state set X_0

- 1 $\Gamma = \text{partition}(X, p)$;
 - /* Search the grid set that satisfies local safety. */
 - 2 $G_1 = \text{constructOneStepGraph}()$;
 - 3 $\Gamma_S, G_K = \text{calculateLocalSafety}()$;
 - /* Search the grid set that satisfies inductiveness. */
 - 4 $\Gamma_I = \text{calculateInductivenessSet}()$;
 - 5 return Γ_I ;
-

also be customized). After fed with the input, the tool works as follows (shown in Algorithm 1). The safe state set X is first uniformly partitioned into small grids $\Gamma = \{v_1, v_2, \dots, v_{p^d}\}$, where $X = v_1 \cup v_2 \cup \dots \cup v_{p^d}$, $v_i \cap v_j = \emptyset$ ($\forall i \neq j$), d is the dimension of the state space, and p is the number of partitions in each dimension (Line 1 in Algorithm 1). The tool then tries to find the grids that satisfy the local safety. It first invokes a reachability graph constructor to build a one-step reachability graph G_1 to describe how the system evolves in one sampling step (Line 2). Then, a dynamic programming (DP) based approach finds the largest set $\Gamma_S = \{v_{s_1}, v_{s_2}, \dots, v_{s_n}\}$ from which the system will not go out of the safe region. The K -step reachability graph G_K is also built in the DP process based on G_1 (Line 3). After that, the tool searches the largest subset Γ_I of Γ_S that satisfies the inductiveness by using a reverse search algorithm (Line 4). The algorithm outputs Γ_I as the target set X_0 (Line 5).

The key functions of the tool are the reachability graph constructor, DP-based local safety set search, and reverse inductiveness set search. In the following sections, we introduce these three functions in detail.

2.1 Reachability Graph Construction

Integration in dynamic system equations is often the most time-consuming part to trace the variation of the states. In this function, we use Flow* to get a valid

Algorithm 2: Construct one-step graph: **constructOneStepGraph()****Data:** Dynamic system f , grid set Γ , the control law π , sampling period δ **Result:** Directed graph $G_1(\Gamma, E_1)$ /* Initialize the edge set E_1 of G_1 . */1 $E_1 \leftarrow \emptyset$;2 **for** $v \in \Gamma$ **do** /* Consider deadline miss ($e = 1$)/meet ($e = 0$) respectively. */3 **for** $e \in \{0, 1\}$ **do** /* Compute one step reachable set $R_1(v)$ from v . */4 $R_1(v) = \text{Flow}^*(v, \delta, e)$; /* v is unsafe and no edge is added if $X^c \cap R_1(v) \neq \emptyset$. */5 **if** $X^c \cap R_1(v) \neq \emptyset$ **then** **Conitnue**; /* Add an edge pointing v' from v if $v' \cap R_1(v) \neq \emptyset$. */6 **for** $v' \in \Gamma$ **do**7 **if** $v' \cap R_1(v) \neq \emptyset$ **then** $E_1 \leftarrow E_1 \cup \{(v, e, v')\}$;8 **return** $G_1(\Gamma, E_1)$;

overapproximation of reachable set (represented as flowpipes) starting from every grid after a sampling period δ . Given a positive integer n , the graph constructed by the reachability set after n sampling period, $n \cdot \delta$, is called a n -step graph G_n . Since the reachability for all the grids in any sampling step is independent under our grid assumption, we first build G_1 and then reuse G_1 to construct G_K later without redundant computation of reachable set.

One-step graph is built with Algorithm 2. We consider deadline miss and deadline meet separately, corresponding to two categories of edges (Line 3). For a grid v , if the one-step reachable set $R_1(v)$ intersects with unsafe state X^c , then it is considered as an unsafe grid and we let its reachable grid be $\emptyset$. Otherwise, if $R_1(v)$ intersects with another grid v' under the deadline miss/meet event e , then we add a directed edge (v, e, v') from v' to v with label e . The number of outgoing edges for each grid node v is bounded by p^d . Assuming that the complexity of Flow^* to compute flowpipes for its internal clock ϵ is $O(1)$, we can get the overall time complexity as $O(|\Gamma| \cdot p^d \cdot \delta/\epsilon)$.

K -step graph G_K is built for finding the grid set that satisfies local safety and inductiveness. To avoid redundant computation on reachable set, we construct G_K based on G_1 by traversing K -length paths, as the bi-product of local safety set searching procedure.

2.2 DP-Based Local Safety Set Search

We propose a bottom-up dynamic programming for considering all the possible paths, utilizing the overlapping subproblems property (Algorithm 3). The reachable grid set at step K that is derived from a grid v at step $k \leq K$ with respect to the number of deadline misses $n \leq m$ can be defined as $\text{DP}(v, n, k)$. To be consistent with Algorithm 2, this set is empty if and only if it does not satisfy the local safety. We need to derive $\text{DP}(v, 0, 0)$. Initially, the zero-step reachability is

Algorithm 3: Search grid set for local safety: **calculateLocalSafety()****Data:** Directed graph $G_1(\Gamma, E_1)$, weakly-hard constraint (m, K) **Result:** Grid set Γ_S , directed graph $G_K(\Gamma, E_K)$

```

1 for  $v \in \Gamma$  do
2   for  $n \leftarrow 0$  to  $m$  do
3      $\text{DP}(v, n, K) \leftarrow \{v\};$ 
4 for  $k \leftarrow K - 1$  to  $0$  do
5   for  $v \in \Gamma$  do
6     for  $n \leftarrow 0$  to  $m$  do
7        $\text{isSafe} \leftarrow \text{True};$ 
8       for  $e \in \{0, 1\}$  do
9         if  $n + e \leq m$  then
10           $\text{nextGrids}(v) \leftarrow \{v' \mid (v, e, v') \in E_1\};$ 
11          if  $\text{nextGrids}(v) = \emptyset$  then  $\text{isSafe} \leftarrow \text{False};$  break;
12          for  $v' \in \text{nextGrids}(v)$  do
13             $R(v') \leftarrow \text{DP}(v, n + e, k + 1);$ 
14            if  $R(v) = \emptyset$  then  $\text{isSafe} \leftarrow \text{False};$  break;
15             $\text{DP}(v', n, k) \leftarrow \text{DP}(v', n, k) \cup R(v);$ 
16          if  $\text{isSafe} = \text{false}$  then
17             $\text{DP}(v, n, k) \leftarrow \emptyset;$ 
18  $\Gamma_S \leftarrow \{v \mid \text{DP}(v, 0, 0) \neq \emptyset\};$ 
19  $E_K \leftarrow \{(v, v') \mid v \in \text{DP}(v, 0, 0)\};$ 
20 return  $\Gamma_S, G_K(\Gamma, E_K);$ 

```

straight forward, i.e., $\forall u \in \Gamma, n \in [0, m], \text{DP}(v, n, K) = \{v\}$. The transition is defined as:

$$\forall k \in [0, K - 1] : \text{DP}(v, n, k) = \bigcup_{\forall v', e: (v, e, v') \in E_1, n + e \leq m} \text{DP}(v', n + e, k + 1).$$

If there exists an empty set on the right hand side or there is no outgoing edge from v for any e such that $n + e \leq m$, we let $\text{DP}(v, n, k) = \emptyset$. Finally, we have $\Gamma_S = \{v \mid \text{DP}(v, 0, 0) \neq \emptyset\}$, $E_K = \{(v, v') \mid v' \in \text{DP}(v, 0, 0)\}$.

We used *bitset* to implement the set union which can accelerate 64 times under the 64-bit architecture. The time complexity is $O(|\Gamma|^2 / \text{bits} \cdot p^d \cdot K^2 + |\Gamma|^2)$, where *bits* depends on the running environment. $|\Gamma|^2$ is contributed by G_K .

2.3 Reverse Inductiveness Set Search

To find the grid set $\Gamma_I \subseteq \Gamma_S$ that satisfies inductiveness, we propose a reverse search algorithm Algorithm 4. Basically, instead of directly searching Γ_I , we try to obtain Γ_I by removing any grid v within Γ_S , from which there exists a path reaching $\Gamma_U = \Gamma - \Gamma_S$. Specifically, Algorithm 4 starts with initializing $\Gamma_U = \Gamma - \Gamma_S$ (line 1). The Γ_U iteratively absorbs the grid v that can reach Γ_U in K sampling periods, until a fixed point is reached (line 2–3). Finally $\Gamma_I = \Gamma - \Gamma_U$ is the largest set that satisfies inductiveness. It is implemented as a breadth first search (BFS) on the reversed graph of G_K , and the time complexity is $O(|\Gamma|^2)$.

Algorithm 4: Search grid set for inductiveness: **calculateInductiveness-Set()**

Data: Directed graph $G_K(\Gamma, E_K)$, Grid set Γ_S

Result: Grid set Γ_I

```

1  $\Gamma_U \leftarrow \Gamma - \Gamma_S$ ;
2 while  $\exists(v, v') \in E_K$  such that  $v \notin \Gamma_U, v' \in \Gamma_U$  do
3   |  $\Gamma_U \leftarrow \Gamma_U \cup \{v\}$ ;
4  $\Gamma_I = \Gamma - \Gamma_U$ ;
5 return  $\Gamma_I$ ;

```

3 Example Usage

Example 1. Consider the following linear control system from [17]:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & -0.1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + u, \quad \text{where } u = \begin{bmatrix} 0 & 0 \\ -0.375 & -1.15 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

$\delta = 0.2$ and $step_size = 0.01$. The initial state set is $x_1 \in [-1, 1]$ and $x_2 \in [-1, 1]$. The safe state set is $x_1 \in [-3, 3]$ and $x_2 \in [-3, 3]$. Following the input format shown in Listing 1.1. Thus, we prepare the model file as Listing 1.2.

```

1 <state_dim> <input_dim> <grid_count>
2 <state_var_names> <input_var_names>
3 <state_ode.1>
4 ...
5 <state_ode.state_dim>
6 <input_equa.1>
7 ...
8 <input_equa.input_dim>
9 <period> <step_size>
10 <m> <k>
11 <safe_state.1>
12 ...
13 <safe_state.state_dim>
14 <initial_state.1>
15 ...
16 <initial_state.state_dim>

```

Listing 1.1. Input format

```

1 2 1 50
2 x1 x2 u
3 x2
4 -0.1 * x2 + u
5 -0.375 * x1 - 1.15 * x2
6 0.2 0.01
7 2 5
8 -3 3
9 -3 3
10 -1 1
11 -1 1

```

Listing 1.2. example/model1.txt

Then, we run our program with the model file.

```
1 ./saw example/model1.txt
```

To further ease the use of our tool, we also pre-compiled our tool for x86_64 linux environment. In such environment, users do not need to compile our tool and can directly invoke **saw_linux_x86_64** instead of **saw** (which is only available after manually compiling the tool).

```
1 ./saw_linux_x86_64 example/model1.txt
```

The program output is shown in Listing 1.3. Line 6 shows the number of edges of G_1 . Lines 8–10 provide the information of G_K , including the number of edges and nodes. Line 12 prints the safe initial set X_0 . Our tool then determines whether the given initial set is safe by checking if it is the subset of X_0 .

```

1 [Info] Parsing model.
2 [Info] Building FLOW* configuration.
3 [Info] Building grids.
4 [Info] Building one-step graph.
5     Process: 100.00%
6 [Success] Number of edges: 19354
7 [Info] Building K-step graph.
8 [Success] Start Region Size: 1908
9     End Region: 1208
10    Number of Edges: 102436
11 [Info] Finding the largest closed subgraph.
12 [Success] Safe Initial Region Size: 1622
13 [Info] Calculating area.
14     Initial state region: 4.000000
15     Grids Intersection: 4.000000
16     Result: safe

```

Listing 1.3. Verification result**Table 1.** Benchmark setting. ODE denotes the ordinary differential equation of the example, π denotes the control law, and δ is the discrete control stepsize.

#	ODE	π	δ	Safe state set	(m, K)
1	$\dot{x}_1 = x_2$ $\dot{x}_2 = -0.1x_2 + u$	$u = -0.375x_1 - 1.15x_2$	0.2	$x_1 \in [-3.0, 3.0]$ $x_2 \in [-3.0, 3.0]$	(2, 5)
2	$\dot{x}_1 = -2x_1 + u_1$ $\dot{x}_2 = -0.9x_2 + u_2$	$u_1 = -x_1$ $u_2 = -x_1 - x_2$	0.3	$x_1 \in [-6.0, 6.0]$ $x_2 \in [-6.0, 6.0]$	(1, 10)
3	$\dot{x}_1 = x_2 + u$ $\dot{x}_2 = -2x_1 - 0.1x_2 + u$	$u = x_1$	1.6	$x_1 \in [-3.0, 3.0]$ $x_2 \in [-3.0, 3.0]$	(2, 10)
4	$\dot{x} = x^2 - x^3 + u$	$u = -2x$	0.6	$x \in [-4.0, 4.0]$	(2, 100)
5	$\dot{x} = 0.2x + 0.03x^2 + u$	$u = -0.3x^3$	1.6	$x \in [-2.0, 2.0]$	(1, 5)
6	$\dot{x}_1 = x_2 - x_1^3 + x_1^2$ $\dot{x}_2 = u$	$u = -1.22x_1 - 0.57x_2$ $-0.129x_2^3$	0.1	$x_1 \in [-5.0, 5.0]$ $x_2 \in [-5.0, 5.0]$	(2, 15)

4 Experiments

We implemented a prototype of SAW that is integrated with Flow*. In this section, we first compare our tool with the one-dimension abstraction [17], on the full benchmarks from [17] (#1–#4) and also additional examples with no guarantee on exponential stability from related works (#5 and #6) [24]. Table 1 shows the benchmark settings, including the (m, K) constraint set for each benchmark. Then, we show how different parameter settings affect the verification results of our tool. All our experiments were run on a desktop, with 6-core 3.60 GHz Intel Core i7.

4.1 Comparison with One-Dimension Abstraction

Table 2 shows the experimental results. It is worth noting that the one-dimension abstraction cannot find the safe initial set in most cases from [17]. In fact, it only

Table 2. Experimental results. ExpParam denotes the parameters of the exponential stability, where “N/A” means that either the system is not exponentially stable or the parameters are not available. Initial state set denotes the set that needs to be verified. The last two columns denote the verification results of the one-dimension abstraction [17] and SAW, respectively. “—” means that no safe initial set X_0 is found by the tool. p represents the partition number for each dimension in SAW. Time (in seconds) represents the execution time of SAW.

#	ExpParam	Initial state set	One-dimension abstraction	SAW		
			Result	p	Result	Time
1	$\alpha = 1.8,$ $\lambda = 0.4$	$x_1 \in [-1.0, 1.0]$ $x_2 \in [-1.0, 1.0]$	—	50	Yes	72.913
2	$\alpha = 1.1,$ $\lambda = 1.8$	$x_1 \in [-6.0, 6.0]$ $x_2 \in [-6.0, 6.0]$	No ($X_0 : x_1^2 + x_2^2 \leq 1.947^2$)	30	Yes	10.360
3	$\alpha = 2,$ $\lambda = 0.37$	$x_1 \in [-1.0, 2.0]$ $x_2 \in [-1.0, 1.0]$	—	100	Yes	183.30
4	$\alpha = 1.4,$ $\lambda = 1$	$x \in [-4.0, 4.0]$	—	30	Yes	80.613
5	N/A	$x \in [-1.56, 1.32]$	—	100	Yes	4.713
6	N/A	$x_1 \in [-5.0, 5.0]$ $x_2 \in [-5.0, 5.0]$	—	50	Yes	750.77

works effectively for a limited set of (m, K) , e.g., when no consecutive deadline misses is allowed. For general (m, K) constraints, one-dimension abstraction performs much worse due to the over-conservation. Furthermore, we can see that, without exponential stability, one-dimension abstraction based approach is not applicable for the benchmarks #5 and #6. Note that for benchmark #2, one-dimension abstraction obtains a non-empty safe initial set X_0 , which however, does not contain the given initial state set. Thus we use “No” instead of “—” to represent this result. Conversely, for every example, our tool computes a feasible X_0 that contains the initial state set (showing the initial state set is safe), which we denote as “Yes”.

4.2 Impact of (m, K) , Granularity, and Stepsize

(m, K) . We take benchmark #1 (Example 1 in Sect. 3) as an example and run our tool under different (m, K) values. Figures 3a, 3b, 3c demonstrate that, for this example, the size of local safety region Γ_S shrinks when K gets larger. The size of inductiveness region Γ_I grows in contrast. Γ_S becomes the same as Γ_I when K gets larger, in which case m is the primary parameter that influences the size of Γ_I .

Granularity. We take benchmark #3 as an example, and run our tool with different partition granularities. The results (Figs. 3d, 3e, 3f) show that Γ_I grows

Fig. 3. Results under different (m, K) values (3a, 3b, 3c) and different granularities (3d, 3e, 3f). The green solid region is Γ_I . The slashed region is Γ_S . The blue rectangle is the initial state set that needs to be verified. (Color figure online)

when p gets larger. The choice of p has significant impact on the result (e.g., the user-defined initial state set cannot be verified when $p = 15$).

Stepsize. We take benchmark #5 as an example, and run our tool with different stepsizes of Flow*. With the same granularity $p = 100$, we get the safe initial state set $\Gamma_I = [-1.56, 1.32]$ when $step_size = 0.1$, but Γ_I is empty when $step_size = 0.3$. The computation times are 4.713s and 1.835s, respectively. Thus, we can see that there is a trade-off between the computational efficiency and the accuracy.

5 Conclusion

In this paper, we present a new tool SAW to compute a tight estimation of safe initial set for infinite-time safety verification of general nonlinear weakly-hard systems. The tool first discretizes the safe state set into grids. By constructing a reachability graph for the grids based on existing tools, the tool leverages graph theory and dynamic programming technique to compute the safe initial

set. We demonstrate that our tool can significantly outperform the state-of-the-art one-dimension abstraction approach, and analyze how different constraints and parameters may affect the results of our tool. Future work includes further speedup of the reachability graph construction via parallel computing.

References

1. Ahrendts, L., Quinton, S., Boroske, T., Ernst, R.: Verifying weakly-hard real-time properties of traffic streams in switched networks. In: Altmeyer, S. (ed.) 30th Euromicro Conference on Real-Time Systems (ECRTS 2018). Leibniz International Proceedings in Informatics (LIPIcs), vol. 106, pp. 15:1–15:22. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl (2018). <https://doi.org/10.4230/LIPIcs.ECRTS.2018.15>. <http://drops.dagstuhl.de/opus/volltexte/2018/8987>
2. Bernat, G., Burns, A., Liamsi, A.: Weakly hard real-time systems. *IEEE Trans. Comput.* **50**(4), 308–321 (2001). <https://doi.org/10.1109/12.919277>
3. Bund, T., Slomka, F.: Controller/platform co-design of networked control systems based on density functions. In: ACM SIGBED International Workshop on Design, Modeling, and Evaluation of Cyber-Physical Systems, pp. 11–14. ACM (2014)
4. 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). https://doi.org/10.1007/978-3-642-39799-8_18
5. Choi, H., Kim, H., Zhu, Q.: Job-class-level fixed priority scheduling of weakly-hard real-time systems. In: IEEE Real-Time Technology and Applications Symposium (RTAS) (2019)
6. Chwa, H.S., Shin, K.G., Lee, J.: Closing the gap between stability and schedulability: a new task model for cyber-physical systems. In: IEEE Real-Time Technology and Applications Symposium (RTAS) (2018)
7. Duggirala, P.S., Viswanathan, M.: Analyzing real time linear control systems using software verification. In: RTSS, pp. 216–226. IEEE (2015)
8. Frehse, G., Hamann, A., Quinton, S., Woehrl, M.: Formal analysis of timing effects on closed-loop properties of control software. In: 2014 IEEE Real-Time Systems Symposium, pp. 53–62, December 2014. <https://doi.org/10.1109/RTSS.2014.28>
9. Frehse, G., et al.: SpaceEx: scalable verification of hybrid systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 379–395. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_30
10. Gaid, M.B., Simon, D., Sename, O.: A design methodology for weakly-hard real-time control. *IFAC Proc.* Vol. **41**(2), 10258–10264 (2008). <https://doi.org/10.3182/20080706-5-KR-1001.01736>. <http://www.sciencedirect.com/science/article/pii/S1474667016406129>, 17th IFAC World Congress
11. Hamdaoui, M., Ramanathan, P.: A dynamic priority assignment technique for streams with (m, k)-firm deadlines. *IEEE Trans. Comput.* **44**(12), 1443–1451 (1995)
12. Hammadeh, Z.A.H., Ernst, R., Quinton, S., Henia, R., Rioux, L.: Bounding deadline misses in weakly-hard real-time systems with task dependencies. In: Design, Automation Test in Europe Conference Exhibition (DATE), pp. 584–589, March 2017. <https://doi.org/10.23919/DATE.2017.7927054>

13. Hammadeh, Z.A.H., Quinton, S., Ernst, R.: Extending typical worst-case analysis using response-time dependencies to bound deadline misses. In: Proceedings of the 14th International Conference on Embedded Software, EMSOFT 2014, pp. 10:1–10:10. ACM, New York (2014). <https://doi.org/10.1145/2656045.2656059>. <http://doi.acm.org/10.1145/2656045.2656059>
14. Hammadeh, Z.A.H., Quinton, S., Panunzio, M., Henia, R., Rioux, L., Ernst, R.: Budgeting under-specified tasks for weakly-hard real-time systems. In: Bertogna, M. (ed.) 29th Euromicro Conference on Real-Time Systems (ECRTS 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 76, pp. 17:1–17:22. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl (2017). <https://doi.org/10.4230/LIPIcs.ECRTS.2017.17>. <http://drops.dagstuhl.de/opus/volltexte/2017/7163>
15. Huang, C., Chen, X., Lin, W., Yang, Z., Li, X.: Probabilistic safety verification of stochastic hybrid systems using barrier certificates. *TECS* **16**(5s), 186 (2017)
16. Huang, C., Wardega, K., Li, W., Zhu, Q.: Exploring weakly-hard paradigm for networked systems. In: Workshop on Design Automation for CPS and IoT (DESTION 2019) (2019)
17. Huang, C., Li, W., Zhu, Q.: Formal verification of weakly-hard systems. In: The 22nd ACM International Conference on Hybrid Systems: Computation and Control (HSCC) (2019)
18. Huang, C., Xu, S., Wang, Z., Lan, S., Li, W., Zhu, Q.: Opportunistic intermittent control with safety guarantees for autonomous systems. In: Design Automation Conference (DAC) (2020)
19. Li, J., Song, Y., Simonot-Lion, F.: Providing real-time applications with graceful degradation of QoS and fault tolerance according to (m, k) -firm model. *IEEE Trans. Industr. Inf.* **2**(2), 112–119 (2006)
20. Liang, H., Wang, Z., Roy, D., Dey, S., Chakraborty, S., Zhu, Q.: Security-driven codesign with weakly-hard constraints for real-time embedded systems. In: 37th IEEE International Conference on Computer Design (ICCD 2019) (2019)
21. Lin, C., Zheng, B., Zhu, Q., Sangiovanni-Vincentelli, A.: Security-aware design methodology and optimization for automotive systems. *ACM Trans. Des. Autom. Electron. Syst. (TODAES)* **21**(1), 18:1–18:26 (2015). <https://doi.org/10.1145/2803174>. <http://doi.acm.org/10.1145/2803174>
22. Marti, P., Camacho, A., Velasco, M., Gaid, M.E.M.B.: Runtime allocation of optional control jobs to a set of CAN-based networked control systems. *IEEE Trans. Industr. Inf.* **6**(4), 503–520 (2010). <https://doi.org/10.1109/TII.2010.2072961>
23. Pazzaglia, P., Pannocchi, L., Biondi, A., Natale, M.D.: Beyond the weakly hard model: measuring the performance cost of deadline misses. In: Altmeyer, S. (ed.) 30th Euromicro Conference on Real-Time Systems (ECRTS 2018). Leibniz International Proceedings in Informatics (LIPIcs), vol. 106, pp. 10:1–10:22. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl (2018). <https://doi.org/10.4230/LIPIcs.ECRTS.2018.10>. <http://drops.dagstuhl.de/opus/volltexte/2018/8993>
24. Prajna, S., Parrilo, P.A., Rantzer, A.: Nonlinear control synthesis by convex optimization. *IEEE Trans. Autom. Control* **49**(2), 310–314 (2004)
25. Quinton, S., Hanke, M., Ernst, R.: Formal analysis of sporadic overload in real-time systems. In: Proceedings of the Conference on Design, Automation and Test in Europe, DATE 2012, EDA Consortium, San Jose, CA, USA, pp. 515–520 (2012). <http://dl.acm.org/citation.cfm?id=2492708.2492836>
26. Ramanathan, P.: Overload management in real-time control applications using (m, k) -firm guarantee. *IEEE Trans. Parallel Distrib. Syst.* **10**(6), 549–559 (1999). <https://doi.org/10.1109/71.774906>

27. Soudbakhsh, D., Phan, L.T., Annaswamy, A.M., Sokolsky, O.: Co-design of arbitrated network control systems with overrun strategies. *IEEE Trans. Control Netw. Syst.* **5**(1), 128–141 (2016)
28. Sun, Y., Natale, M.D.: Weakly hard schedulability analysis for fixed priority scheduling of periodic real-time tasks. *ACM Trans. Embed. Comput. Syst. (TECS)* **16**(5s), 171 (2017)
29. Wardega, K., Li, W.: Application-aware scheduling of networked applications over the low-power wireless bus. In: *Design, Automation and Test in Europe Conference (DATE)*, March 2020
30. Xu, W., Hammadeh, Z.A.H., Kröller, A., Ernst, R., Quinton, S.: Improved deadline miss models for real-time systems using typical worst-case analysis. In: *2015 27th Euromicro Conference on Real-Time Systems*, pp. 247–256, July 2015. <https://doi.org/10.1109/ECRTS.2015.29>
31. Yang, Z., Huang, C., Chen, X., Lin, W., Liu, Z.: A Linear programming relaxation based approach for generating barrier certificates of hybrid systems. In: Fitzgerald, J., Heitmeyer, C., Gnesi, S., Philippou, A. (eds.) *FM 2016*. LNCS, vol. 9995, pp. 721–738. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48989-6_44
32. Zhu, Q., Sangiovanni-Vincentelli, A.: Codesign methodologies and tools for cyber-physical systems. *Proc. IEEE* **106**(9), 1484–1500 (2018). <https://doi.org/10.1109/JPROC.2018.2864271>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.