Verified Switched Control System Design using Real-Time Hybrid Systems Reachability

Stanley Bak, Taylor Johnson, Marco Caccamo, Lui Sha Air Force Research Lab – Information Directorate – Rome, NY

DISTRIBUTION A. Approved for public release; Distribution unlimited. (Approval AFRL PA #88ABW-2014-2661, 02 JUNE 2014)





- Include computational components interacting with the physical-world
- Mistakes can have real-world consequences!
- Ideally we would verify the system, but it may be too complicated for direct verification



Fault-Tolerant Power Distribution



Autonomous Cars



Air Traffic Control







• Run-Time Assurance (RTA) Design

• RTA using Real-Time Reachability







• Run-Time Assurance (RTA) Design

• RTA using Real-Time Reachability







- Sandbox untrusted controllers
- Lots of variants
- Key challenge is decision module





• Safe design is easy!

- The challenge is **conservatism**. The 'best' switching logic:
 - Predicts next state using the current command
 - Checks if the safety controller recovers





 Using simulations, we can grid the state space, then check which states are recoverable*



* M. Aiello, J. Berryman, J. Grohs, and J. Schierman, "Run-time assurance for advanced flight-critical control systems," in Proceedings of the American Institute of Aeronautics and Astronautics Guidance, Navigation, and Control Conference, ser. AIAA '10, 2010.





- When do we stop the simulation?
- What if the real state is between simulation points?
- How accurate are the simulations?
- Most problematic: How well does this scale?







Most problematic: How does this scale?

- 100 partitions per dimension,
 11 dimensions = 100^11 = 10^22 points,
 1 µs per simulation = **317 million years**
- Large online storage required
 - Lookup table?
 - Linear bounds*?

* S. Bak, "Industrial application of the System-Level Simplex Architecture for real-time embedded system safety," Master's Thesis, University of Illinois at Urbana-Champaign, 2009.





How many partitions per dimension if we want the runtime to be:

- 1 hour ~7 partitions (~51 degrees)
- 1 day ~ 10 partitions (~36 degrees)
- 1 year ~ 17 partitions (~20 degrees)





- Instead of simulation, we can use more formal reasoning based on hybrid-systems reachability computation
- This reasons about sets of states, accounting for method inaccuracies by over-approximation







- The 'best' switching logic can be defined in terms of reachability: REACH(REACH $_{\delta}(x, CC), SC) \cap U = \emptyset$
 - -or- $x \notin \text{BACKREACH}^*_{\leq \delta}(\text{BACKREACH}^*(U, \text{SC}), \text{CC}')$
- Drawbacks:
 - Achievable Accuracy
 - Online Representation





- Flow* is a tool which computes reachability for systems with nonlinear dynamics
 - Uses Taylor Models for representation

- Example: 9 dimensional biological model*
 Order: 5, Step size: 0.001, Steps: 10
 - Output: 3 MB, ~300 KB per step







 For linear (and linearized) systems, you can find the largest ellipsoid inside the recoverable region by solving a linear matrix inequality (LMI)*

Linear Time-Invariant Control System: x' = Ax + Bu

- Input: Matrices A, B, linear system constraints
- Output: Gain matrix K, Potential matrix P, where if you use u=Kx then x^TPx is decreasing and all constraints are satisfied if x^TPx < 1

* D. Seto and L. Sha, "A case study on analytical analysis of the inverted pendulum real-time control system," CMU/ SEI, Tech. Rep., 1999.

LMI-design for RTA





LMI-design for RTA (2)









1. Ellipsoid guarantees input will not saturate, which is pessimistic

2. Ellipsoid may trim out recoverable states because of its shape restriction











- From before: The 'best' switching logic:
 - Predicts next state using the current command
 - Checks if the safety controller recovers

• For the offline LMI approach, you consider all possible commands and how much state space can be covered in one control iteration, and then create a 'buffer'















• Run-Time Assurance (RTA) Design

• RTA using Real-Time Reachability





- Don't determine the switching set offline, do it online!
 - No large enumeration
 - Not limited to ellipsoid shape
 - No complex state representation

- How do we do it?
 - Use aspects of both LMI-based and reachability-based Simplex design





- We use the forward-time definition of a switching set $\operatorname{REACH}(\operatorname{REACH}_{\delta}(x,\operatorname{CC}),\operatorname{SC}) \cap U = \emptyset$
- However, we don't need infinite time reachability; we only need to get back into the LMI ellipsoid because

$$REACH(\mathcal{R}, SC) \cap U = \emptyset$$
$$\mathcal{R} = \{x | x^T P x < 1\}$$





• **Key Idea**: allow the system to leave the safe ellipsoid, as long as we can guarantee (1) no constraints are violated when this happens, and (2) the state is guaranteed to go back into the ellipsoid







- This requires reachability at runtime
 - Tools aren't meant for this...

- Let's make one!
 - Based on mixed-face lifting
 - Quick computation is more important than long-term error control
 - Assumes piecewise dynamics, with bounded derivatives, and a user-provided
 DerivativeBounds function





• For our algorithm, the user must provide a function that bounds the derivative for each direction in an arbitrary box







Tracked States













Real-Time Reachability Algorithm



Maximum advancement of each face after desiredStep time







Construct neighborhood for each face based on the widths







- Next step: resample derivatives within neighborhoods
 - Reconstruct if either (1) neighborhood flips from inwards-facing to outwards-facing, or (2) estimated width doubles in size







Compute time until one of the faces can leave its neighborhood





Advance time and repeat

Tracked States

Next Tracked States





- Construction time is bounded (finite derivative can only double a finite number of times)
- Time to advance per step is at least half of desired step size (number of steps is bounded)
- Typically, O(n) to advance time, but poor long-term error control due to wrapping effect
- After finishing the computation for the desired reachtime, <u>cut desired step size in half and repeat</u> until deadline exhausted

Iterative Improvement









• How well does this compare to LMI-based reachability and simulation (optimal)?

 Use 4-D saturated inverted pendulum (same system as the one used for LMI-based design)







• Projection where position = velocity = 0



Evaluation (2)



• Projection where θ =0.19, ω =0.18







- Partition state space and count the points in each region
- "Best" improvement (based on simulations) is 247%

Runtime (ms)	LMI	RealTime	Sim	Unrecov	Improve
5	5473	6180	1376	37596	213%
20	5473	6948	608	37596	227%
40	5473	7059	497	37596	229%
50	5473	7108	448	37596	230%
75	5473	7133	423	37596	230%
100	5473	7216	340	37596	232%
200	5473	7286	270	37596	233%
500	5473	7338	218	37596	234%
1000	5473	7382	174	37596	235%







- "A Unified Run-Time Assurance Scheme using Real-Time Reachability"
- Run-Time Assurance Design
 - Simulation
 - Reachability
 - Linear Matrix Inequality (LMI)
- RTA using a <u>Unified Approach</u>
 - LMI ellipsoid + online reachability
 - Real-time reachability is feasible, even for small runtimes
 - Unified approach greatly expands the complex controller region (227% on the saturated inverted pendulum system)