College of Engineering
UtahStateUniversity

Enhancing Property-Directed Reachability for Chemical Reaction Networks

Landon Taylor and Zhen Zhang, Electrical and Computer Engineering Department

Project Goals

Chemical Reaction Networks (CRNs) are a formal language for modeling chemical kinetics, modeled as continuous-time Markov chains (CTMCs). CRNs include rare events that challenge state-of-the-art tools [1]. While probabilistic extensions of PDR are promising [2, 3], these tools are designed for *discrete-time* Markov models. Further, these tools *verify* a *user-provided* upper bound on the probability of failure, requiring expert knowledge of the model or guess-and-check methods.

We tackle three main **challenges** for CRN verification:

1. CRN analysis requires enumerating *many* traces, especially for rare events;
2. CRN verification requires full numerical analysis and an explicit state space; and
3. Bounding variables in infinite-state CRNs requires deep insights into the model.

We propose the following two **solutions** to these problems:

1. An automated method to bound a state space by analyzing PDR frames; and
2. An alternative to explicit state space enumeration: an expansion to PDR that symbolically estimates the probability of reaching a rare event $\neg P$ in a CRN.

Property-Directed Reachability (PDR)

Property-Directed Reachability (PDR) is a highly-scalable symbolic verification technique designed for reachability analysis of Boolean systems [4, 5]. PDR has since been expanded and revised to verify systems with quantifier free formulae over bitvectors [6] and Markov decision processes [3].

- PDR generates a trace from initial state I to a *target state* violating safety property P or providing an invariant as a proof of P.
- PDR analyzes a model using relatively inductive frames. Frame F_k contains reachable states from F_{k-1} in one step.
- In traditional PDR, execution terminates when the whole reachable state space satisfies P or when PDR cannot exclude a target state from the reachable area.
- This work leverages PDR's ability to find traces to $\neg P$ to obtain the probability of a CRN behavior of interest.

A Simple PDR Example

Frame F_0 shows that only I is reachable in zero steps. In F_1, unreachable regions are blocked (shown in gray). F_1 proves $\neg P$ is unreachable from I in 1 step. Then F_2 is created, and unreachable regions from F_1 are blocked. More states are reachable, but $\neg P$ is still blocked. After creating F_3, we cannot block every $\neg P$-state. The red counterexample trace is returned to show P fails.

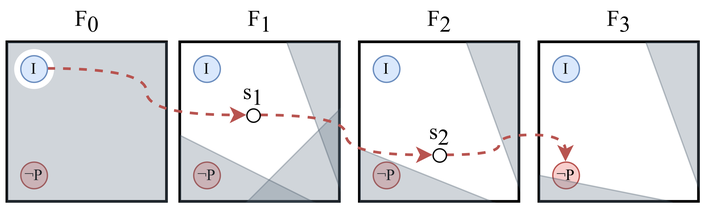

This inductive checking procedure makes PDR an extremely efficient tool to find counterexample traces and check safety properties in deterministic models. We propose expanding its functionality to probabilistic models.

Using PDR for Variable Bounding

CRNs often have large or infinite state spaces. In our experience, knowing how to bound variables in a complex model becomes problematic for users. For example, we are often required to set extremely high bounds for variables to avoid excluding an important part of a state space.

This causes a dilemma for a CRN verifier:

- If the bound is **too high**, state explosion causes time and memory problems.
- If the bound is **too low**, we may exclude valuable states from our state space.

We suggest PDR can provide valuable insights to set reasonable bounds. It already finds unreachable regions of a state space, so it follows that these regions can be used to generate variable bounds.

Backward Analysis

Consider the following CRN modification:

- Consider $\neg I$ as the safety property
- Consider $\neg P$ as the initial state formula
- Reverse all transitions

If PDR checks this backward model, we obtain a backward-reachable region. States in this region can reach a target state within n steps, regardless of the initial state. If the forward-reachable region is Region A, and the backward-reachable region is Region B, the region $A \cap B$ contains **all** the traces that start at I and end at $\neg P$. This gives us a tighter state space, allowing for accurate species bounds and lower verification effort.

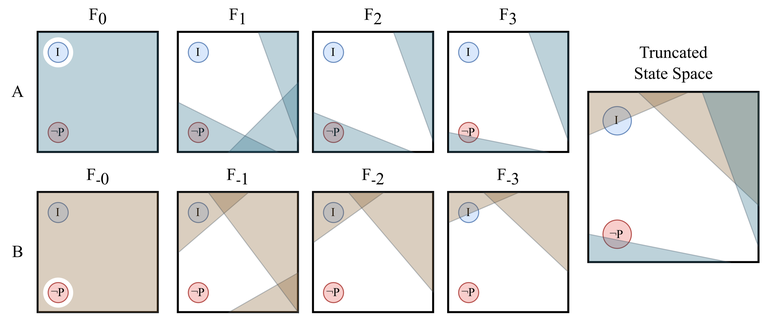

The probability returned by probabilistic model checking using these bounds is a lower bound on the true probability because the model is only guaranteed to include traces up to a finite number of steps.

PDR Adaptations for Variable Bounding

We propose three adaptations to PDR to optimize it for CRN variable bounding:

1. Use information from the CRN model to suggest smart bounds to PDR. For example, if x can increase 2, a suggested bound at step k is $x \leq 2k$.
2. After one n-step counterexample is found, frame F_n contains *every* n-step trace. CRNs are highly concurrent, so allowing an additional few steps can significantly increase the probability of a CRN reaching $\neg P$. Variable bounds can be derived from frame F_{n+k}, providing all traces of up to $n + k$ steps.
3. Bounds can be further tightened by evaluating backward reachability.

Layered PDR For Probability Estimates

Our goal is to obtain the probability of a CRN to reach a rare event of interest within a time bound. We propose a layered approach to PDR for incremental probability estimation:

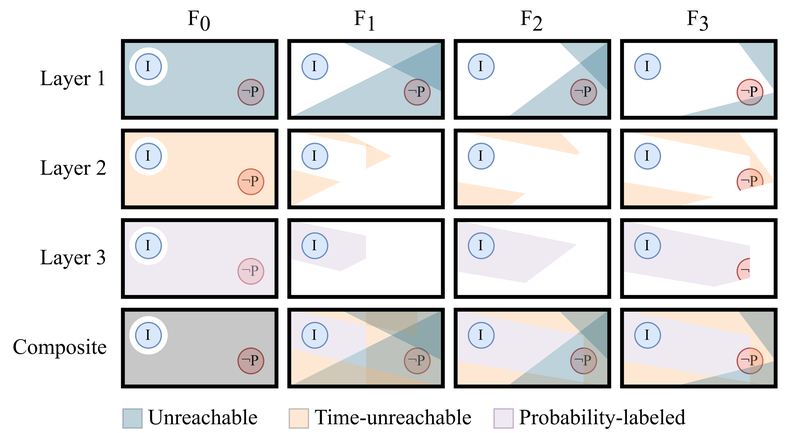

Three layers combine to form the modified PDR:

- **Layer 1: Reachability.** Layer 1 is our adapted implementation of PDR.
- **Layer 2: Time-Bounding.** Layer 2 analyzes CRN time-bounded reachability information. Using PDR's inductive checking, this layer accumulates mean dwell times (average time spent in CTMC states) and blocks regions that take too long to reach a target state.
- **Layer 3: Probability-Labeled Regions** This layer groups states into regions of states with similar properties and labels each region with a probability upper and lower bound. The final probability upper and lower bounds can then be computed inductively starting at regions satisfying $\neg P$ in the final frame.

By superimposing these layers, we find an *estimate* of the probability. We suggest this method will be able to return relatively tight bounds around the true probability.

References & Acknowledgement

This work was supported by the National Science Foundation under Grant No. 1856733. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

[1] B. Israelsen, L. Taylor, and Z. Zhang, "Efficient trace generation for rare-event analysis in chemical reaction networks," in *Model Checking Software* (G. Caltais and C. Schilling, eds.), (Cham), pp. 83–102, Springer Nature Switzerland, 2023.

[2] E. Polgreen, M. Brain, M. Fraenzle, and A. Abate, "Verifying Reachability Properties in Markov Chains via Incremental Induction," Sept. 2019.

[3] K. Batz, S. Junges, B. L. Kaminski, J.-P. Katoen, C. Matheja, and P. Schröer, "PrIC3: Property Directed Reachability for MDPs," in *Computer Aided Verification* (S. K. Lahiri and C. Wang, eds.), vol. 12225, pp. 512–538, Springer International Publishing, 2020.

[4] A. R. Bradley, "SAT-Based Model Checking without Unrolling," in *Verification, Model Checking, and Abstract Interpretation* (R. Jhala and D. Schmidt, eds.), Lecture Notes in Computer Science, pp. 70–87, Springer, 2011.

[5] N. Een, A. Mishchenko, and R. Brayton, "Efficient implementation of property directed reachability," in *Proceedings of the International Conference on Formal Methods in Computer-Aided Design*, FMCAD '11, (Austin, Texas), pp. 125–134, FMCAD Inc, Oct. 2011.

[6] T. Welp and A. Kuehlmann, "QF_BV Model Checking with Property Directed Reachability," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013*, pp. 791–796, IEEE Conference Publications, 2013.