

Probabilistic Program Equivalence for NetKAT

Steffen Smolka
Cornell University

David Kahn
Cornell University

Praveen Kumar
Cornell University

Nate Foster
Cornell University

Dexter Kozen
Cornell University

Alexandra Silva
University College London

Abstract

We study the problem of deciding program equivalence in the context of Probabilistic NetKAT, a formal language for reasoning about the behavior of packet-switched networks. We show that the problem is decidable for the history-free fragment of the language, and discuss a path toward a decision procedure for the full language. The main challenge lies in reasoning about iteration, which we address by a reduction to finite-state absorbing Markov chains. We also describe an OCaml prototype that we have used to reason about probabilistic network programs.

1 Motivation

NetKAT [1] is a language based on Kleene algebra (KA) and Kleene algebra with tests (KAT) that can be used to program networks and reason about their properties. It comes with a rich theory (a denotational semantics, a sound and complete axiomatization, both language and automata models) and sophisticated tools (a fully-automatic decision procedure, an efficient compiler). Probabilistic NetKAT [4] extends the language with a probabilistic choice operator and a semantics based on Markov kernels, allowing it to model features such as faulty links, randomized routing algorithms, and uncertainty about input packets.

Why study program equivalence? Many network properties can be naturally and conveniently phrased as questions about program equivalence in NetKAT. For example, NetKAT equivalence has been used to reason about essential properties such as waypointing, reachability, isolation, and loop freedom, as well as for the validation and verification of compiler transformations. This work aims to extend this approach to the probabilistic setting.

Why is it challenging? Because ProbNetKAT has an iteration operator, it is possible to write programs that generate continuous distributions over the uncountable space of packet history sets. This makes reasoning about convergence non-trivial, and raises the issue of representing infinitary objects in an implementation. Prior work developed a domain-theoretic semantics that provides notions of approximation and continuity, which can be used to reason about programs using distributions with finite-support [4]. However, it left program equivalence as an open problem. We settle this question positively for the history-free fragment of the language, in which distributions have finite-support

but iteration may still converge only in the limit. Computing the analytical value of these limits precisely is the main challenge we address.

Why not apply standard results? Although there is a lot of related work (e.g., probabilistic languages based on regular operators [3], automata [5], and model checkers [2]), ProbNetKAT’s semantics seems to be sufficiently different that previous results do not apply since programs in the history-free fragment do not denote sequences in any obvious sense and correspond to one-state automata.

Our Approach. Our decision procedure for the history-free fragment of ProbNetKAT follows a general approach: we map programs to canonical representations for which checking equivalence is straightforward. Specifically, we define a *big-step* semantics that interprets each program as a finite stochastic matrix—equivalently, a Markov chain that transitions from input to output in a single step. Equivalence is trivially decidable on this representation, but some care is needed to compute the matrix in the case of iteration—intuitively, it needs to capture the result of an infinite stochastic process. We address this by embedding the system in a second Markov chain with a larger state space that models iteration in the spirit of a *small-step* semantics. This chain can be transformed to an absorbing Markov chain, which admits a closed form analytic solution using elementary matrix operations that represents the limit of the iteration. We have proved the soundness of this approach.

2 Results

Preliminaries. A history is a non-empty sequence of packets that encodes the trajectory of a single packet through the network: it can be thought of as a log of the packet’s activity. On a semantic level, ProbNetKAT programs p denote functions of type $\llbracket p \rrbracket \in 2^H \rightarrow \mathcal{D}(2^H)$ that map sets of input packet histories to distributions on sets of output packet histories. On a syntactic level, the language consists of a Boolean algebra (which enables using predicates to classify incoming packets based on their contents); the well-known regular operators $+$, \cdot , and $*$ (which enables specifying regular forwarding paths); a probabilistic choice operator \oplus ; and a modification primitive $f \leftarrow n$ (which allows rewriting the f -field of incoming packets to n). The primitive `dup` is a “logging” command that records the current state of the packet in the history. It is important to note that the notion

- [2] M. Kwiatkowska, G. Norman, and D. Parker. 2011. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *Proc. 23rd International Conference on Computer Aided Verification (CAV'11) (LNCS)*, G. Gopalakrishnan and S. Qadeer (Eds.), Vol. 6806. Springer, 585–591. https://doi.org/10.1007/978-3-642-22110-1_47
- [3] A. K. McIver, E. Cohen, C. Morgan, and C. Gonzalia. 2008. Using Probabilistic Kleene Algebra pKA for Protocol Verification. *J. Logic and Algebraic Programming* 76, 1 (2008), 90–111.
- [4] Steffen Smolka, Praveen Kumar, Nate Foster, Dexter Kozen, and Alexandra Silva. 2017. Cantor Meets Scott: Semantic Foundations for Probabilistic Networks. In *POPL 2017*. <https://doi.org/10.1145/3009837.3009843>
- [5] Wen-Guey Tzeng. 1992. A Polynomial-Time Algorithm for the Equivalence of Probabilistic Automata. *SIAM J. Comput.* 21, 2 (1992), 216–227. <https://doi.org/10.1137/0221017>