

Probabilistic Verification of BGP Convergence

Soroush Haeri
Simon Fraser University
Vancouver, British Columbia, Canada
Email: shaeri@sfu.ca

Dario Krešić
University of Zagreb
Varaždin, Croatia
Email: dario.kresic@foi.hr

Ljiljana Trajković
Simon Fraser University
Vancouver, British Columbia, Canada
Email: ljilja@sfu.ca

Abstract—The *Border Gateway Protocol* (BGP) is the de facto Internet routing protocol. Various aspects of the BGP protocol have been analyzed using mathematical and experimental approaches. Formal verification of BGP specification validates whether or not a specific set of requirements is satisfied. In recent years, the probabilistic behavior of BGP has been explored. The size of routing tables has been modeled as a stochastic process that changes over time according to some probability distribution function. Hence, the verification of BGP may also be probabilistic in nature due to its randomized behavior. In this paper, we present a probabilistic model checking approach to analyze BGP convergence properties that may be employed to automate the BGP convergence analysis.

I. INTRODUCTION

Border Gateway Protocol (BGP) [1] has been used widely as the main inter *Autonomous System* (AS) Internet routing protocol. The extensive growth of the Internet over the past decade resulted in routing tables with over 300,000 entries [2]. This makes the analysis of BGP convergence properties more difficult and calls for employment of automated techniques.

An AS selects its preferred routes based on its routing policy and the best routes that have been advertised by its neighboring ASes. Local AS policies play an important role in preferred route selection because the BGP allows policy-based decisions to override distance metrics. Local routing policies are usually defined based on a limited knowledge of other ASes policies and network topology and, hence, may be inconsistent. This may cause that a set of ASes exchange route information messages infinitely and do not converge to a set of stable routes.

The BGP convergence has been widely studied. Griffin and Wilfong introduced a timeless formal BGP model called *Stable Path Problem* (SPP) to analyze BGP convergence properties [3]. Obradovic extended the model with real-time information to observe the evolution of the protocol over time [4]. Viswanathan et al., introduced to SPP the concept of probability and examined the convergence properties of various topologies using a probabilistic model [5].

Model checking, introduced and developed by Clarke and Emerson [6], is an automated technique to formally verify the correctness of a finite-state system. Correctness is defined with respect to a set of required properties to be fulfilled by the system. In the context of model checking, these properties are known as specifications. For example, *safety* and *liveness* are two important specifications for communication protocols. Safety may be described as an event or set of events that

should never occur while liveness is a desired property that should eventually happen [7]. In recent years, model checking has been extended to verification of probabilistic systems [8]–[11].

Input to model checking process is a variant of finite-state systems and required specifications expressed in terms of temporal logic. If the system satisfies a specification, the process returns "yes". Otherwise, it returns a counterexample path falsifying the specification. In contrast, stochastic model checking expresses the likelihood of an event to occur. By adding the probability operator to temporal logic one may express the probability of occurrence of an event in a given time interval rather than whether or not such event occurs [12].

In this paper, we propose a probabilistic model checking approach to analyze safety of a BGP policy and the BGP convergence time using *Probabilistic Computation Tree Logic* (PCTL) [9], [12]. We use the probabilistic BGP model introduced by Viswanathan et al., [5] as the underlying formalism for the analysis. The proposed approach may be used to automate the analysis of BGP convergence.

In Section II, we briefly introduce the probabilistic BGP model. In Section III, we define simple PCTL expressions that prove safety of an arbitrary BGP policy and its convergence time. Conclusions and future work are given in Section IV.

II. GLOBAL BGP EXECUTION MODEL

Assume that node 0 is the single destination for all other nodes. Let $G = (V, E)$ be an undirected graph and $\mathcal{P} = \{\mathcal{P}^v \mid v \in V - \{0\}\}$ be a set of permitted paths, where \mathcal{P}^v is a set of paths such that each nonempty path in \mathcal{P}^v is a simple path from v to the destination node. Let $\Lambda = \{\lambda^v \mid v \in V - \{0\}\}$ be a ranking function on $\mathcal{P}^v \rightarrow \mathbb{N}$, where \mathbb{N} denotes set of natural numbers. An instance of *stable path problem* (SPP) [3] is defined as a triple $S = (G, \mathcal{P}, \Lambda)$ provided that empty path ϵ is permitted with the lowest rank ($\lambda^v(\epsilon) = 0$) and that *strictness* and *simplicity* are also satisfied. Definitions and a detailed model description are given in [3].

Viswanathan et al., [5] described the global BGP execution model based on SPP as an input-output automaton with $Q(S)$ being the set of states of the global automaton describing path assignments. Let π be a mapping function assigning each node $u \in V - \{0\}$ to a permitted path $\pi(u) \in \mathcal{P}^u$. We assume for every node $u \in V - \{0\}$ that the path assignment initially maps the empty path ϵ to u ($\pi(u) = \epsilon$). Inputs to the automaton are events of the form $\{advertise_u \mid u \in U\}$

for some $U \subseteq V - \{0\}$. Transitions of such automaton from an arbitrary state π to π' based on an input event $e = \{advertise_u \mid u \in U\}$ (i.e., $\pi \xrightarrow{e} \pi'$) is defined by $\pi'(v) = \pi(v)$ for any $v \notin U$ and $\pi'(u) = Next_u(\pi)(u)$ for any $u \in U$. Let $choices(u, \pi)$ denote the candidate paths available at node u under an arbitrary path assignment π . Next path chosen by router u as its preferred path is defined as $Next_u(\pi)(u) = \max(u, choices(u, \pi))$.

The transition matrix $\mathbf{T}(S)$ corresponds to an instance of SPP $S = ((V, E), \mathcal{P}, \Delta)$ and is of dimensions $|Q(S)| \times |Q(S)|$, where $\mathbf{T}(S)_{ij} = \{U \mid \pi_i \xrightarrow{\{advertise_u \mid u \in U\}} \pi_j\}$.

Let $\mathbf{p} = (p_1, \dots, p_n)$ be an activation probability vector with $n = |V| - 1$, where each p_i represents the probability that node i receives the event $advertise_i$ and recompute its routes. The BGP speakers have no global coordinations and, hence, we may assume that $advertise_i$ events are independent.

The transition matrix $\mathbf{T}(S)$ may evolve to a stochastic transition matrix $\mathbf{T}'(S)$ by casting operator $P(\cdot)$ on every element of $\mathbf{T}(S)$ ($\mathbf{T}'(S)_{ij} = P(\mathbf{T}(S)_{ij})$). Operator $P(\cdot)$ is defined as

$$P(U) = \left(\prod_{i \in U} p_i \right) \prod_{j \notin U} (1 - p_j), \quad (1)$$

where p_i is the i^{th} element of the activation vector \mathbf{p} and U is a set of nodes that recompute their best path after occurrence of an event $\{advertise_u \mid u \in U\}$. For an arbitrary set γ , which is a subset of power set of $V - \{0\}$,

$$P(\gamma) = \sum_{U \in \gamma} p(U). \quad (2)$$

III. PCTL FOR VERIFICATION OF BGP CONVERGENCE

Contrary to solving a system of linear equations to calculate the convergence time [5], we propose a PCTL rewarding process that may be employed to explore convergence time of various systems. We provide a PCTL expression that verifies the safety of a configuration. The model checking approach is simple and various systems may be tested using a unique set of PCTL expressions. It also eliminates the need to derive equations for each specific network and it enables considering larger network topologies.

$\mathbf{T}'(S)$ is a stochastic matrix that represents the transition matrix of a *discrete time Markov chain* (DTMC). The properties of a system may be verified by probabilistic model checking approach. *Safety* and *convergence time* are the most important properties for BGP convergence.

A. Safety

Any instance of SPP is safe with respect to an initial state π_0 if and only if for an activation probability vector \mathbf{p} there is no cyclic state. The system does not enter any state infinitely often unless it is the absorbing state. The expression

$$P_{>=1}[\mathbf{GF} \pi \rightarrow \mathbf{FG} \pi] \quad \forall \pi \in Q(S), \quad (3)$$

where \mathbf{G} and \mathbf{F} are linear time logic operators and \rightarrow denotes implication, verifies safety of a global BGP execution instance.

B. Convergence time

To calculate the convergence time we define a state reward function $\rho(\pi)$ as

$$\rho(\pi) = 1 \quad \forall \pi \in Q(S). \quad (4)$$

Each time a system enters a state from its set of acceptable states, it receives a reward equal to one. This rewarding process enables counting the number of states visited before system reaches the absorbing state implying that the system converges and δ denotes a unique absorbing state. The number of transitions made until convergence is reached at state δ may be expressed as

$$\mathcal{R}_{=?}[\mathbf{F} \delta]. \quad (5)$$

The reward operator is $\mathcal{R}_{\sim r}[\mathbf{F} \Phi]$, where $\sim \in \{<, \leq, >, \geq\}$ and Φ denotes an arbitrary state. The bound r has to be satisfied before Φ is reached. PCTL reward checking algorithm first calculates the total reward and then compares it with the reward bound r . Total reward achieved before a state satisfying Φ , denoted by $\mathcal{R}_{=?}[\mathbf{F} \Phi]$, may be calculated using the PCTL model checking algorithm [12].

IV. CONCLUSION

In this paper, we introduced an approach to analyze the convergence properties of BGP by employing model checking techniques. The proposed approach will be tested on BGP implementations with various routing policies and network topologies.

REFERENCES

- [1] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," IETF, RFC 1771, March 1995.
- [2] G. Huston. (2011, July) Analysis of BGP routing table dynamics. [Online]. Available: <http://bgp.potaroo.net/>.
- [3] T. G. Griffin and G. Wilfong, "An analysis of BGP convergence properties," in *Proc. Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, New York, NY, USA, September 1999, pp. 277–288.
- [4] D. Obradovic, "Real-time model and convergence time of BGP," in *Proc. INFOCOM*, New York, US, June 2002, pp. 893–901.
- [5] R. Viswanathan, K. K. Sabnani, R. J. Holt, and A. N. Netravali, "Expected convergence properties of BGP," *Computer Networks*, vol. 55, no. 8, pp. 1957–1981, June 2011.
- [6] E. M. Clarke and E. A. Emerson, "Design and synthesis of synchronization skeletons using branching-time temporal logic," in *Logic of Programs Workshop*, London, UK, 1982, pp. 52–71.
- [7] G. Holzmann, *Design and Validation of Computer Protocols*. New Jersey, US: Prentice Hall, 1991.
- [8] R. Alur, C. Courcoubetis, and D. Dill, "Model-checking for probabilistic real-time systems (extended abstract)," in *Proc. 18th International Colloquium on Automata, Languages and Programming*, Madrid, Spain, July 1991, pp. 115–126.
- [9] H. Hansson and B. Jonsson, "A logic for reasoning about time and reliability," *Formal Aspects of Computing*, vol. 6, no. 5, pp. 512–535, September 1994.
- [10] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton, "Model-checking continuous-time Markov chains," *ACM Transactions on Computational Logic*, vol. 1, no. 1, pp. 162–170, July 2000.
- [11] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, "Model-checking algorithms for continuous-time Markov chains," *IEEE Transactions on Software Engineering*, vol. 29, no. 6, pp. 524–541, June 2003.
- [12] M. Kwiatkowska, G. Norman, and D. Parker, "Stochastic model checking," *Formal Methods for Performance Evaluation*, vol. 4486, pp. 220–270, 2007.