

Reduction-based Analysis of BGP Systems with BGPVerif

Anduo Wang* Alexander J.T. Gurney* Xianglong Han* Jinyan Cao*
Carolyn Talcott† Boon Thau Loo* Andre Scedrov*
*University of Pennsylvania †SRI International
{anduo,jinyan,hanxiang,agurney,boonloo}@cis.upenn.edu, clt@csl.sri.com,
scedrov@math.upenn.edu

ABSTRACT

Today’s inter-domain routing protocol, the Border Gateway Protocol (BGP), is increasingly complicated and fragile due to policy misconfiguration by individual autonomous systems (ASes). Existing configuration analysis techniques are either manual and tedious, or do not scale beyond a small number of nodes due to the state explosion problem. To aid the diagnosis of misconfigurations in real-world large BGP systems, this paper presents *BGPVerif*, a reduction based analysis toolkit. The key idea is to reduce BGP system size prior to analysis while preserving crucial correctness properties. *BGPVerif* consists of two components, *NetReducer* that simplifies BGP configurations, and *NetAnalyzer* that automatically detects routing oscillation. *BGPVerif* accepts a wide range of BGP configuration inputs ranging from real-world traces (Rocketfuel network topologies), randomly generated BGP networks (GT-ITM), Cisco configuration guidelines, as well as arbitrary user-defined networks. *BGPVerif* illustrates the applicability, efficiency, and benefits of the reduction technique, it also introduces an infrastructure that enables networking researchers to interact with advanced formal method tool.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols

General Terms

Verification, Management

Keywords

Border gateway protocol, reduction, formal analysis

1. INTRODUCTION

The Internet today runs on a complex routing protocol called the *Border Gateway Protocol* or *BGP* for short. BGP enables autonomous systems (ASes) worldwide to achieve global connectivity, subject to each system’s local policy (what paths are allowed, and the route preference used to select best paths). The convergence behavior of the global Internet depends on how each ISP configures its policy. Prior work has shown that policy misconfigurations can lead to route oscillation and slow convergence [2].

While our understanding of BGP safety has improved dramatically in the past decade, there is a lack of automated tools that can

aid in the safety analysis process. Analysis today either consists of manual proofs and counter-examples on small instances, or custom simulators that are unable to explore the full protocol state space.

In this demonstration, we present *BGPVerif*, a formal analysis toolkit that enables networking researchers to study and analyze large BGP systems in a sound and automatic fashion. *BGPVerif* consists of two components. *NetAnalyzer* [10] is an automated analysis toolkit based on the Maude rewriting logic engine [4]. In *NetAnalyzer*, the BGP system is encoded as a transition system driven by rewriting rules. Maude’s high-performance rewriting engine is used to analyze BGP instances automatically, by simulating execution runs, as well as exhaustively exploring all execution runs for possible divergence.

While automated, *NetAnalyzer* is susceptible to the state explosion problem as the number of nodes increases. The *NetReducer* component applies a novel reduction theory [9, 8] for reducing BGP network size while preserving convergence properties. In combination, *NetAnalyzer* and *NetReducer* enable *BGPVerif* to perform safety analysis efficiently, often completing the analysis on large networks that would otherwise not be possible to study within reasonable time.

2. OVERVIEW

As shown in Figure 1, *BGPVerif* takes as input the configuration of a BGP system, which includes the network topology and routing policies for all routers. Policies include routing filters and preferences based on path attributes. *BGPVerif* converts this configuration to an *Extended Path Digraph* (EPD) [8] instance, an internal representation used by *BGPVerif* to perform the reduction. The EPD structure includes an explicit representation of both the network topology and the preference-induced dependencies among paths. EPD extends prior models [1, 6], and provides a basis for reducing instances prior to analysis [9].

Given the EPD representation of the BGP system, the *NetReducer* component first performs a reduction to output the reduced EPD. The basic idea of reduction is straightforward: to simplify the BGP system of interest, one repeatedly merges all pairs of nodes where the merger does not affect the convergence behavior of the entire network. There are two rules for identifying such node pairs, and each of them only requires checking of local policy details. (1) *Duplicate* reduction merges two nodes whose routing configurations agree upon what paths and route preferences they learned from their neighbors. (2) *Complementary* reduction merges two nodes whose neighbors agree upon the paths and preferences learned from them.

In our demonstration prototype, reduction can perform reduction in two modes: (1) In *batch* mode, *BGPVerif* repeatedly merges all reducible nodes and outputs the final reduced BGP system in the

form of a EPD; (2) Alternatively, in the *interactive* mode, BGPVerif allows the user to select which nodes to try to reduce.




Figure 1: BGPVerif System Overview.

Finally, given a BGP configuration’s EPD (either before or after reduction), the NetAnalyzer component uses the Maude analyzer [10, 9, 3] to classify the configuration as either convergent or oscillatory, and the oscillation as either transient or permanent.

3. DETAILS OF DEMONSTRATION

BGPVerif provides a graphical user interface for the user to specify BGP configuration inputs, and invoke reduction or analysis. The interface embeds a visualizer that demonstrates the dynamic reduction process as a transformation of the network topology augmented with routing policies. In the back-end, BGPVerif includes the reduction engine NetReducer and the analysis engine NetAnalyzer based on Maude. NetReducer is responsible for computing the internal EPD representation, and performing the actual reduction, whereas NetAnalyzer extends Maude with additional library support for detecting route oscillations in BGP systems.

Specify Configuration.

In the BGPVerif GUI, the user inputs the topology and policy of the BGP configuration of interest, or selects a configuration already provided. The user can edit these pre-specified configurations as well. The built-in network topologies include real world Rocketfuel data [7], and randomly generated networks using GT-ITM [5], which is controlled by the input parameters. The built-in routing policies include standard iBGP configurations such as full mesh, route reflection and confederations [11].

Once the BGP system configuration is loaded, BGPVerif generates the internal EPD representation by running BGP in a special mode that deduces all possible paths and their relative preferences. The user can now choose *batch* or *interactive* reduce commands from NetReducer, or the *analyze* command of NetAnalyzer.

NetReducer: Simplify BGP Systems.

In batch mode, NetReducer will try all combinations of nodes to see which are reducible. This makes the batch mode useful as an automated step prior to analysis. In the fine-grained interactive mode, the user can specify either one or two nodes that she wants to merge. The NetReducer can check whether a given pair of nodes is indeed reducible; or, given only a single node, can search for a second node to form a reducible pair. The interactive mode provides insights into convergence behavior even without a following analysis operation.

Consider the scenario where a network operation would like to expand her BGP system by adding new routers, without changing the convergence behavior. Rather than analyzing the two BGP systems before and after migration separately, she can instead incrementally grow the BGP network size while preserving convergence throughout. At each step, the operator adds a node, and by interactive reduction, she can decide if it can be merged with any of the

existing nodes. She can also add a new router in a chosen place by deliberately making it duplicate or complementary with an existing node.

NetAnalyzer: Detect Route Oscillations.

NetAnalyzer provides an interactive graphical interface for networking researchers to make use of the formal analyses implemented in Maude [10, 9, 3]. It allows the user to analyze BGP configurations using reduction techniques that accelerate analysis by orders of magnitude, enabling automated analysis of configurations that could not previously be handled.

NetAnalyzer utilizes Maude’s built-in exhaustive search ability and adds library support for detecting BGP oscillations. Our library allows us to exhaustively search all execution runs to detect route oscillation. On termination, NetAnalyzer can correctly classify a BGP configuration as either safe (no oscillations in any protocol runs), transient oscillation (route oscillation happens in some execution traces), or permanent oscillation (convergence is impossible).

4. ACKNOWLEDGMENT

This research is funded in part by NSF grants (CCF-0820208, CNS-0830949, CNS-0845552, CNS-1040672, TC-0905607 and CPS-0932397), AFOSR grant FA9550-08-1-0352, and ONR grant N00014-11-1-0555.

5. REFERENCES

- [1] GRIFFIN, T. G., AND WILFONG, G. An analysis of BGP convergence properties. In *SIGCOMM* (1999).
- [2] LABOVITZ, C., MALAN, G., AND JAHANIAN, F. Internet Routing Instability. *TON* (1998).
- [3] MASON, I. A., AND TALCOTT, C. L. Iop: The interoperability platform & imaude: An interactive extension of maude. *Electron. Notes Theor. Comput. Sci.* 117 (Jan. 2005), 315–333.
- [4] MAUDE. <http://maude.cs.uiuc.edu/>.
- [5] MODELING TOPOLOGY OF LARGE INTERNETWORKS. <http://www.cc.gatech.edu/projects/gtitm/>.
- [6] SOBRINHO, J. Network routing with path vector protocols: theory and applications. In *SIGCOMM* (2003).
- [7] SPRING, N., MAHAJAN, R., AND WETHERALL, D. Measuring ISP topologies with Rocketfuel. In *ACM SIGCOMM* (2002).
- [8] WANG, A., GURNEY, A. J. T., TALCOTT, C., LOO, B. T., AND SCEDROV, A. A calculus of policy-based routing systems. In *PODC* (2012). Brief announcement.
- [9] WANG, A., TALCOTT, C., GURNEY, A. J. T., LOO, B. T., AND SCEDROV, A. Reduction-based formal analysis of BGP instances. 18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2012).
- [10] WANG, A., TALCOTT, C., JIA, L., LOO, B. T., AND SCEDROV, A. Analyzing BGP instances in Maude. In *13th IFIP International Conference on Formal Methods for Open Object-based Distributed Systems. 31th IFIP International Conference on FORmal TEchniques for Networked and Distributed Systems. (FMOODS-FORTE 2011)*.
- [11] ZHANG, R., AND BARTELL, M. *BGP Design and Implementation*. Cisco Press, 2003.