

Analyzing BGP Instances in Maude

Anduo Wang¹ Carolyn Talcott² Limin Jia³ Boon Thau Loo¹ Andre Scedrov¹

¹ University of Pennsylvania

² SRI International

³ Carnegie-Mellon University

{anduo, boonloo}@cis.upenn.edu clt@csl.sri.com
liminjia@cmu.edu scedrov@math.upenn.edu

Abstract. Analyzing Border Gateway Protocol (BGP) instances is a crucial step in the design and implementation of safe BGP systems. Today, the analysis is a manual and tedious process. Researchers study the instances by manually constructing execution sequences, hoping to either identify an oscillation or show that the instance is safe by exhaustively examining all possible sequences. We propose to automate the analysis by using Maude, a tool based on rewriting logic. We have developed a library specifying a generalized path vector protocol, and methods to instantiate the library with customized routing policies. Protocols can be analyzed automatically by Maude, once users provide specifications of the network topology and routing policies. Using our Maude library, protocols or policies can be easily specified and checked for problems. To validate our approach, we performed safety analysis of well-known BGP instances and actual routing configurations.

1 Introduction

The Internet today runs on a complex routing protocol called the *Border Gateway Protocol* or *BGP* in short. BGP enables Internet-service providers (ISP) world-wide to exchange reachability information to destinations over the Internet, and simultaneously, each ISP acts as an autonomous system that imposes its own import and export routing policies on route advertisements exchanged among neighboring ISPs.

Over the past few years, there has been a growing consensus on the complexity and fragility of BGP routing. Even when the basic routing protocol converges, conflicting policy decisions among different ISPs have led to route oscillation and slow convergence. Several empirical studies [11] have shown that there are prolonged periods in which the Internet cannot reliably route data packets to specific destinations due to routing errors induced by BGP. In response, the networking community has proposed several alternative Internet architectures [18] and policy constraints (or “safety guidelines”) that guarantee protocol convergence if universally adopted [8,6,16].

One of the key requirements for designing new routing architectures and policy guidelines is the ability to study BGP network instances. These instances can come in the form of small topology configurations (called “gadgets”), which serve as examples of safe systems, or counterexamples showing the lack of convergence. They can also come from actual internal router (iBGP) and border gateway (eBGP) router configurations. Today, researchers and network operators analyze these instances by manually examining execution sequences. This process is tedious and error-prone.

The main contribution of this paper is that we present an automated tool for analyzing BGP instances, and thus relieve researchers and network operators of manual analysis. Our tool uses Maude [4], a language and tool based on rewriting logic. We encode in Maude the BGP protocol as a transition system driven by rewriting rules. Consequently, we can use the high-performance rewriting engine provided by Maude to analyze BGP instances automatically. Our tool can simulate execution runs, as well as exhaustively explore all execution runs for possible divergence.

More specifically, we developed a set of Maude libraries specifying a generalized path vector protocol that is common to all BGP instances. The generalized path vector protocol utilizes a set of unspecified routing policy functions. These unspecified functions serve as the interface for specific routing policies which are formalized as *Stable Path Problems (SPP)* [10]. To use our library, users only need to input the network topology and customize routing policies functions in the form of SPP. We illustrate the use of our tool by analyzing various BGP instances.

2 Background

2.1 BGP

BGP assumes a network model in which routers are grouped into various Autonomous Systems (AS) administrated by Internet Service Provider (ISP). Individual AS exchanges route advertisements with neighboring ASes using the *path-vector* protocol. Upon receiving a route advertisement, a BGP node may choose to accept or ignore the advertisement based on its *import policy*. If the route is accepted, the node stores the route as a possible candidate. Each node selects among all candidate routes the best route to each destination based on its local route rankings. Once a best route is selected, the node advertises it to its neighbors. A BGP node may choose to export only selected routes to its neighboring ASes based on its *export policy*.

BGP systems come in two flavors: external BGP (eBGP), which establishes routes between ASes; and internal BGP (iBGP), which distributes routes within an AS. At the AS-level, a BGP system can be viewed as a network of AS nodes running eBGP. Each AS is represented by one single *router* node in the network (its internal structure ignored), and its *network state* includes its neighbors (ASes), selected best path and a routing table. Route advertisements constitute the *routing messages* exchanged among them. Within an AS, a BGP system can be viewed as a network of two kinds of network nodes running iBGP: gateway routers and internal routers whose network states are similar to eBGP routers. iBGP allows internal routers to learn external routes (to destination outside the AS) from gateway routers.

We model both eBGP and iBGP systems as network systems with two components: routing dynamics and routing policies. Routing dynamics specifies how routers exchange routing messages, and how they update their network states accordingly. Routing policies are part of the static configuration of each router, by which the ISP operator expresses his local traffic interests and influences route decisions.

In our library, we adopt the use of Stable Paths Problems (SPP) [10] as the formal model of routing policies. An instance of the SPP $S = (G, o, P, A)$, is a graph G , a spe-

cific destination node o ⁴, the set of permitted (usable) paths P available for each node to reach o , and the ranking functions for each node A . For each node v , λ^v is its ranking function, mapping its routes to natural numbers (ranks), and P^v are its permitted paths, the set of available paths to reach o . A path assignment is a function π that maps each network node $v \in V$ to a path $\pi(v) \in P^v$. A path assignment is *stable* if each node u selects a path $\pi(u)$ which is (1) the highest ranked path among its permitted paths, and (2) is consistent with the path chosen by the next-hop node. Consistency requires if $\pi(u) = (uv)P$ then for the next-hop node v , we must have $\pi(v) = P$. A solution to the SPP is a stable path assignment.

In this paper, we are interested in analyzing BGP convergence (safety) property in the SPP formalism. A BGP system converges and is said to be safe, if it produces stable routing tables, given any sequences of routing message exchanges. We can study BGP convergence by analyzing its SPP representation: SPP instance for a safe BGP system converges to a solution in all BGP executions. Note that, the existence of a SPP solution does not guarantee convergence.

For example, Figure 1 presents a SPP instance called the Disagree “gadget”. The per-node ranking functions are $\lambda^1([n1\ n2\ n0]) = 1$, $\lambda^1([n1\ n0]) = 2$, $\lambda^2([n2\ n1\ n0]) = 1$, and $\lambda^2([n2\ n0]) = 2$. The permitted paths for each node are listed besides the corresponding node.

The order in which the paths are listed is based on the ranking function: Nodes prefer higher ranked

routes, e.g. node $n1$ prefers route $[n1\ n2\ n0]$ over $[n1\ n0]$. Disagree has two stable path assignment solutions: $([n1\ n2\ n0], [n2\ n0])$ and $([n2\ n1\ n0], [n1\ n0])$. However, Disagree is not guaranteed to converge to any of the stable solutions because there exists a particular execution trace where routes keep oscillating. In particular, in the execution of BGP protocol, if node $n1$ and $n2$ update and exchange routing messages in a synchronized manner, their network states can oscillate between two unstable path assignments $([n1\ n0], [n2\ n0])$ and $([n1\ n2\ n0], [n2\ n1\ n0])$ forever.




Fig. 1: Disagree Gadget.

2.2 Rewriting Logic and Maude

Rewriting logic [13] is a logical formalism that is based on two simple ideas: states of a system can be represented as elements of an algebraic data type, and the behavior of a system can be given by transitions between states described by *rewrite rules*. By algebraic data type, we mean a set whose elements are constructed from atomic elements by application of constructors. Functions on data types are defined by equations that allow one to compute the result of applying the function. A rewrite rule has the form $t \Rightarrow t'$ if c where t and t' are patterns (terms possibly containing variables) and c is a condition (a boolean term). Such a rule applies to a system in state s if t can be matched to a part of s by supplying the right values for the variables, and if the condition c holds when supplied with those values. In this case the rule can be applied by replacing the part of s matching t by t' using the matching values for variables in t' .

⁴ Assuming the Internet is symmetric, we can study its routing behavior by studying routing to a single destination.

Maude [4] is a language and tool based on rewriting logic [12]. Maude provides a high performance rewriting engine featuring matching modulo associativity, commutativity, and identity axioms; and search and model-checking capabilities. Given a specification S of a concurrent system, Maude can execute this specification and allows one to observe possible behaviors of the system. One can also use the search functionality of Maude to check if a state meeting a given condition can be reached during the system’s execution. Furthermore, one can model-check S to check if a temporal property is satisfied, and if not, Maude will produce a counter example. Maude also supports object-oriented specifications that enable the modeling of distributed systems as multiset of objects that are loosely coupled by message passing. As a result, Maude is particularly amenable to the specification and analysis of network routing protocols.

3 A Maude Library for Encoding BGP Protocols

This section presents our Maude library for analyzing BGP instances. This library provides specification of the protocol dynamics that are common to BGP instances, and defines a routing policy template in terms of the Stable Path Problem (SPP) so that network designers can customize it to analyze a specific instance. Our library also provides support for detecting route oscillation.

BGP system	Maude interpretation
Network nodes	(Router) objects
Routing messages	Terms of type <code>Msg</code>
Global Network	Multiset of router objects and terms representing messages
Protocol dynamics	Local rewriting rules
Global network behaviors	Concurrent rewrites using local rules
Route oscillation support	A <code>Logger</code> object recording the histories of route assignments and rewriting rules updating the <code>Logger</code> object

Table 1: Overview and Interpretation of Maude Library

Our library is organized into a hierarchy of Maude modules. Table 1 presents the correspondence between concepts in BGP protocols and the Maude code. We first show how our library represents a single network state of BGP system (Section 3.1). Then we explain how to capture the dynamic behavior of a local BGP router using rewrite rules. In doing so, the global network behaviors can be viewed as concurrent applications of the local rewriting rules (Section 3.2). Finally, we discuss the component in the library that detects route oscillation (Section 3.3).

3.1 Network State

A network state is represented by a multiset of network nodes (routers) and routing messages used by routers to exchange routing information. Each network node is represented by a Maude object, whose attributes consist of its routing table, best path and

neighboring table. We omit the detailed Maude sort definitions, but provide an example encoding of the network node `n1` in Disagree gadget show in Figure 1 as follows.

```
[n1 : router |
  routingTable: (source: n1,dest: n0,pv:(n1 n0),metric: 2),
  bestPath: (source: n1,dest: n0,pv:(n1 n0),metric: 2),
  nb: (mkNeigh(n0,2) mkNeigh(n2,1))]
```

The constructor for a node is `[_:_|_,_,_]_`. The first two elements (`n1:router`) specifies the node's id `n1`, and its object class `router`. The next three elements are the attributes. At a given state, the routing table attribute constructed from `routingTable:..` contains `n1`'s current available routes. Each routing table entry stores the routing information for one particular next-hop. Here, the routing table attribute only contains one entry (`source: n1, dest: n0, pv:(n1 n0), metric: 2`). This route is specified by its source (`source: n1`), destination (`dest: n0`), the path vector that contains the list of nodes along the path (`pv: (n1 n0)`), and the cost of the route (`metric: 2`). This route is also used for the best path attribute, constructed from `bestPath:..`, which contains `n1`'s current best path. The last attribute is the neighbor table, constructed from `nb:..`. To extract a node's local neighbor table from the network topology, we further introduce an operator `mkNeigh`. The first argument of `mkNeigh` is the identifier of the neighboring node, and the second argument the metric associated with the link to that node. Node `n1` has two neighbors, node `n0`, the cost to which is 2 (`mkNeigh(n0,2)`); and node `n2`, the cost to which is 1 (`mkNeigh(n2,1)`).

Besides router objects, the second part of a network state is routing messages in the network. Typically, network nodes exchange routing information by sending each other routing messages carrying newly-learned routes. In our library, a routing message is constructed from `sendPacket(-,-,-,-,-)`. For example, in the Disagree gadget, the initial routing message sent by node `n1` to its neighbor `n2` is represented by message term: `sendPacket(n1,n2,n0,2,n1 n0)`. This message carries `n1`'s routes to destination `n0` with path vector `n1 n0` at cost 2. In general, the first two arguments of `sendPacket(-,-,-,-,-)` denote the sender's identifier (node `n1`), and the receiver's identifier (node `n2`) respectively. The rest of the arguments specify the identifier of the destination (node `n0`), the metric representing the cost of the route (2), and the path vector of the routing path (`n1 n0`).

3.2 Protocol Dynamics

We now show how to specify network system dynamics in Maude. By modeling a BGP system as a concurrent system consisting of router objects (and the routing messages), to specify the global BGP evolution, we only need to specify the local rewrite rules governing the state transition of each BGP router.

A BGP node's dynamics can be captured by various equivalent state transitions. To reduce search space in analysis, we adopt a one-state transition: for each BGP node `N`, when it receives routing messages from a neighbor `S`, `N` computes the new path from the received message, updates `N`'s routing table and re-selects best path accordingly, and finally sends out routing messages carrying its new best path information if a different best path is selected. This state transition is mapped to a single rewrite rule of the following form:

```

rl [route-update] :
  sendPacket(S, N, D, C, PV)
  [ N : router | routingTable: RT, bestPath: Pb, nb: NB ]
=>
if (case 1) then best path re-selects (promotion)
else (if (case 2) then best path remains same
      else (if (case 3) then best path re-selection (withdraw)
            else error processing
            fi) fi) fi.

```

This rule is fired when the left-hand side is matched; that is, when a node N consists of routingTable RT , bestPath Pb , and neighboring table NB receives a route advertisement message from neighbor S . The result of applying the rule is shown on the right-hand side: the routing message is consumed, and attributes of router N are updated. Based on the result of the re-selected bestPath attribute, there are three different cases for N to update its state as specified in the three branches. Next, we explain these three cases.

Best path promotion. In any case, node N needs to first compute the new path based on its neighbor S 's message asserting that S can reach D via a path PV . We define a function `newPath` that takes a routing message and the neighbor table as arguments, and returns the new path by first prepending N to the path announced by S , setting the new path attribute according to the local ranking function `lookUpRank`, and then imposing the import policy by modifying the path metric according to BGP routing policy configuration (`import` function). Here `import` and `lookUpRank` are unspecified routing policy functions. Together with `export` that we will introduce shortly, they constitute our library's specification interface for defining BGP routing policy. To specify a particular BGP instance's routing policy, the user only needs to specify `import`, `lookUpRank` and `export` accordingly.

In the first branch (case 1), this newly computed path is then compared with the current bestPath Pb , if the new one is preferred over the old value Pb , the `bestPath` attribute will be updated to this new path. Furthermore, if the export policy allows, the new best path value will be re-advertised to all of N 's neighbors by sending them routing messages, specified as follows:

```

if getDest(newPath(sendPacket(S,N,D,C,PV),NB))==getDest(Pb) and
  prefer?(newPath(sendPacket(S,N,D,C,PV),NB),Pb)==true
then
([ N : router |
  routingTable: updatedRT(newPath(sendPacket(S,N,D,C,PV),NB),RT),
  bestPath: newPath(sendPacket(S,N,D,C,PV),NB),
  nb: NB ]
multiCast(NB, export(newPath(sendPacket(S,N,D,C,PV),NB))))

```

Here the new state of N is obtained by updating the old routingTable attribute RT (`updateRT` function), and updating bestPath attribute by setting it to the new value of bestPath. The `updateRT` function recursively checks the routing table, and for each next-hop entry, it either inserts the new `newPath(...)` if no available route is presented; or replaces the old less-preferred value with the new `newPath(...)`. To complete the state transition, for all N 's neighbors, routing messages carrying `newPath(...)`

are generated by `multiCast` function. To impose the export routing policy, before sending the new best path, `export` is applied to `newPath(...)` to filter out the routes which are intended to be hidden from neighbors. Similar to `import`, `export` is to be instantiated by the user when analyzing a particular BGP instance. If the export routing policy prohibits `newPath(...)` to be announced, `export` will transform it to `emptyPath`, which `multiCast` transforms to `none`, the empty message set.

Best path remains the same. In the second branch (case 2), a new path `newPath(...)` is computed from the received message as before. However, the new path is no better than the current `bestPath Pb`. But the next-hop node of `newPath(...)` and `Pb` are different, implying that `newPath(...)` is just an alternative path⁵ for `N` to reach the destination. As a result, the current `bestPath` value `Pb` is unchanged, and only the `routingTable` will be updated with this alternative path `newPath(...)`. No routing messages will be generated:

```
if getDest(newPath(sendPacket(S,N,D,C,PV),NB))==getDest(Pb) and
   getNext(newPath(sendPacket(S,N,D,C,PV),NB))!=getNext(Pb) and
   prefer?(Pb,newPath(sendPacket(S,N,D,C,PV),NB))==true
then
[ N : router |
  routingTable: updateRT(newPath(sendPacket(S,N,D,C,PV),NB),RT),
  bestPath: Pb,
  nb: NB ]
```

Best path withdraw. As in the second branch, in case 3, the newly computed path `newPath(...)` is worse than the current `bestPath Pb`, but it is now routed through the same next-hop `S` as current `bestPath Pb`. The fact that `S` now sends a less preferred path indicates that the previous learned better route `Pb` is no longer available at `S`. Therefore, we need to withdraw `Pb` by dropping `Pb` from routing table, re-selected and re-advertised a new best path (cascading withdraw), shown as follows:

```
if getDest(newPath(sendPacket(S,N,D,C,PV),NB))==getDest(Pb) and
   getNext(newPath(sendPacket(S,N,D,C,PV),NB))==getNext(Pb) and
   prefer?(Pb,newPath(sendPacket(S,N,D,C,PV),NB))==true
then
([ N : router |
  routingTable: updateRT(newPath(sendPacket(S,N,D,C,PV),NB),RT),
  newBest(newPath(sendPacket(S,N,D,C,PV),NB),
           updateRT(newPath(sendPacket(S,N,D,C,PV),NB),RT)),
  nb: NB ]
multiCast(NB,export(newBest(newPath(sendPacket(S,N,D,C,PV),NB),
                             updateRT(newPath(sendPacket(S,N,D,C,PV),NB),
                             RT))))
```

Here, `updateRT` replaces (therefore removes) the outdated `Pb` with the less-preferred `newPath(...)`, and `newBest` is the function that re-computes the best path from `newPath(...)` and the remaining paths in routing table. As in case 1, to complete state transition, the newly selected best path is sent to its neighbors by `multiCast(...)`.

⁵ Different next-hop implies the route is learned from a different neighbor.

3.3 Route Oscillation Detection Support

Our library also provides extra definitions to help detect route oscillation. Our method is based on the observation that if route oscillation occurs during network system evolution, there is at least one *path assignment* (at a given state for a BGP system, we define the path assignment to be the collection of best paths currently selected by all network nodes) that is visited twice. Therefore, we use the following simple heuristic: we maintain a record of all path assignments for all visited states in BGP execution, and check for recurring path assignment. Note that a path assignment (best path attribute of `router` object) only constitutes a sub-set of the entire system state (the `router` objects attributes and routing messages), consequently our heuristic based on this partial system state can have false positive problem: our heuristic may report a false route oscillation by catching recurring states in only path assignments, but not the entire system states. Nevertheless, our heuristic is sound and is still helpful in detecting all potential route oscillation: when route oscillation occurs, a recurring path assignment state must occur.

Specifically, in our Maude library, we create a global *logger object* to keep track of the history of path assignments. For each snapshot of the network state, i.e. whenever a network node makes a local state transition and updates its best path attribute, the logger object is synchronized to create a new path assignment entry that corresponds to the updated best path. We then provide a function that checks recurring entries in the list of visited path assignments, which can be used directly in Maude's exhaustive search to detect route oscillation.

Logger object. The global logger is represented by an object `pa` of `Logger` class which has one attribute: At a given state, the attribute contains a list (history) of path assignments, each entry of which contains the snapshot of the network's collection of best paths in a visited state. At the top of the list, the first path assignment entry shows the network's current (latest) best paths. An example logger object for the disagree gadget is the following:

```
{pa : Logger | history: ({[n1 n2 n0] [n2 n0]}
                        {[n1 n2 n0] [n2 n1 n0]}
                        {[n1 n2 n0] [n2 n0]}
                        {[n1 n0] [n2 n0]})}
```

The above logger records four snapshots of the Disagree's best paths. For example, the first path assignment `{[n1 n2 n0] [n2 n0]}` denotes the network latest state where node 1's best path to 0 is `[n1 n2 n0]` and node 2's best path is `[n2 n0]`. And line 4 `{[n2 n0] [n2 n0]}` records Disagree's path assignment at its initial (oldest) state. Note that, this object content actually exhibits route oscillation (line 1 and line 3) described in Section 3.2.

Synchronized logging. To log all path assignment changes, we only need to slightly modify the single rewrite rule for route update, such that whenever the rule is fired to apply local state transition for some node, the global object `pa` is synchronized and its path assignment is updated to reflect changes in the local node's best path attribute, shown as follows:


```

rl [route-update-logging] :
  sendPacket(S, N, D, C, PV)
  [ N : router | routingTable: RT, bestPath: Pb, nb: NB ]
  { pa : Logger | history: HIS }
=>
*** first branch: bestPath re-selects (promotion)
if ... then ...
  { pa : Logger | history:
    historyAppend(updateAt(index(N),
      [getPV(newPath(sendPacket(S,N,D,C,PV),NB)]),
      head(HIS)),HIS)}
else ... fi .

```

On the left-hand side, two objects: a router N and the global logger pa are matched to trigger the transition. As described in 3.2, in the first branch of route update where the node's best path attribute is set to $newPath(\dots)$, the logger pa updates its path assignment attribute as follows: First, it creates a new path assignment entry to record $newPath(\dots)$ by function $updateAt(\dots)$. Then, the new entry $updateAt(\dots)$ is inserted into the list of previous path assignments HIS by function $historyAppend$. Here, the new path assignment entry $updateAt(\dots)$ is computed by updating the latest path assignment entry $head(HIS)$ with $newPath(\dots)$. The rest of branches 2 and 3 are modified similarly.

Route oscillation detection. A network state is now a multiset of router objects, routing messages, and one global logger object. The function $detectCycle$ detects re-occurring path assignments, as follows:

```

eq detectCycle([ N : router | routingTable: RT,
                    bestPath: Pb,nb: NB] cf)
  = detectCycle (cf) .
eq detectCycle(message cf) = detectCycle (cf) .
eq detectCycle({ pa : Logger | history: HIS } cf)
  = containCycle? (HIS) .

```

The first two equations ignore router objects and routing messages in the network state, and the last equation examines logger pa by function $containCycle?$ to check for recurring path assignment entries in HIS . We will revisit the use of $detectCycle$ to search for route oscillation in Section 5.

4 Specifying BGP Instance

Given a BGP instance with its *network topology* and *routing policies*, we show how to specify the instance as a SPP in our library. We discuss examples for both eBGP and iBGP.

4.1 eBGP instance

An eBGP instance can be directly modeled by a SPP instance $S = (G, o, P, \Lambda)$: G, o specifies the instance's *network topology*, and P, Λ specifies the resulting per-node route

ranking function after applying the eBGP instance's *routing policies*. Our library provides Maude definitions for each SPP element.

Network topology. An eBGP instance's initial network state is generated from its network topology, which is represented by a list of network nodes and links. Our library declares two constants `top-Nodes` and `top-BGP` to represent network nodes and links. For example, to specify the topology of the Disagree gadget, the user defines `top-Nodes`, `top-BGP` as follows:

```
eq top-Nodes = n1 n2 .
eq top-BGP = (n1,n0 : 2) (n1,n2 : 1) (n2,n1 : 1) (n2,n0 : 2) .
```

Here, `n0` is the identifier of the destination node (*o*). Each link is associated with its cost. Based on the value of `top-Nodes` and `top-BGP` that are input by the user, our library automatically generates Disagree's initial state by `init-config` function:

```
eq gadget = init-config (top-Nodes, top-BGP) .
```

The resulting `gadget` is a network state which consists of the two network router objects `n1, n2`, the four initial routing messages, and the initial logger `pa`, as shown in Section 5.1. In this initial state, the three attributes of each network node – the routing table and best-path and neighbor tables are computed as follows: `init-config` parses the BGP links in network topology (`top-BGP`), for each link `(ni,nj : M)`, a new routing table entry for `nj` with cost `M` is created, and if `nj == n0`, then set `ni`'s best path to the one-hop direct path `ni n0`, and its routing tables containing this one-hop direct route; otherwise if there is no direct link from `ni` to `n0`, set `ni`'s best path and the routing table to `emptyPath`. Initial routing messages and logger `pa` are computed in a similar manner.

Routing policy. The route ranking function Λ and permitted paths P are the result of applying three BGP policies functions: `import`, `export` and `lookUpRank`. As we have discussed in Section 3, `import`, `export`, `lookUpRank` are three user-defined functions that serve as the specification interface for routing policies.

`import` and `lookUpRank` are used in computing new routing paths from a neighbor's routing message: `import` filters out un-wanted paths, and `lookUpRank` assigns a rank to the remaining permitted paths.

Note that the metric value `lookUpRank (N PV)` assigned by `lookUpRank` also determines the route's preference in route selection. `export` is used to filter out routes the router would like to hide.

As an example, the policy functions for Disagree are defined as follows.

```
eq export (P) = P . eq import (P) = P .
eq lookUpRank (n1 n2 n0) = 1 . eq lookUpRank (n1 n0) = 2 .
eq lookUpRank (n2 n1 n0) = 1 . eq lookUpRank (n2 n0) = 2 .
```

The first line says Disagree does not employ additional import/export policies. Whereas the second and third line asserts that Disagree's two nodes prefers routes through each other: For example the second line encodes node `n1`'s ranking policy that it prefers path `(n1 n2 n0)` (with higher rank 1) through `n2` over the direct path `(n1 n0)` (rank 2).

4.2 iBGP Instance

Our technical report [19] shows our SPP encoding of iBGP instances. The main differences between an iBGP and eBGP instances are: (1) iBGP network topology distinguishes between internal routers and gateway routers. Gateway routers runs eBGP to exchange routing information with (gateway routers of) other ISPs, while simultaneously running iBGP to exchange the external routing information with internal routers in the AS. (2) iBGP routing policy utilizes a separate IGP protocol to select best route. Internal to an AS, the ISP uses his own IGP protocol to compute shortest paths among all routers. The shortest path distance between internal routers and gateway routers are used in iBGP route selection: iBGP policy requires the internal routers to pick route with shortest distance to its gateway router.

As a result, iBGP requires encoding two types of topologies: a *signaling* topology for gateway routers and internal routers to exchange routes within the AS, and a *physical* topology on which the IGP protocol is running. Further, an additional *destination router* denoting the special SPP destination o is added as an external router which is connected with all gateway routers. In our library, we implement and run separately in Maude an IGP protocol (for computing all-pairs shortest paths) and pass the resulting shortest path distances to iBGP protocol.

5 Analysis

To analyze BGP instances, our library allows us to (1) execute the Maude specification to simulate possible execution runs; and (2) exhaustively search all execution runs to detect route oscillation.

5.1 Network Simulation

Network initialization. For any analysis, we need to first generate a BGP instance's initial network state. For a given BGP instance, we have shown how to generate its initial state `gadget` from its network topology and routing policy, as described in section 4. For example, the initial state generated for Disagree is as follows:

```
{pa : Logger | history: {[n1 n0] [n2 n0]}}
[n1 : router | routingTable: (source: n1, dest: n0,
                             pv: (n1 n0), metric: 2),
               bestPath: (source: n1, dest: n0,
                           pv: (n1 n0), metric: 2),
               nb: (mkNeigh(n0,2) mkNeigh(n2,1))]
[n2 : router | ... ]
sendPacket(n1,n0,n0,n2,n1 n0) sendPacket(n1,n2,n0,n2,n1 n0)
sendPacket(n2,n0,n0,n2,n2 n0) sendPacket(n2,n1,n0,n2,n2 n0)
```

This state consists of Disagree's initial logger object `pa` that holds the initial path assignment `[n1 n0] [n2 n0]`, two router objects `n1, n2`, and four initial routing messages.

Execution. Unlike many formal specification paradigms used in static network analysis, Maude specification is executable. To explore *one* possible execution run from a given initial state `gadget`, we can directly use Maude’s `rewrite` and `frewrite` (fair rewriting) commands. For example, we could tell Maude to execute the Disagree gadget with the following command:

```
Maude > frew gadget .
```

This command terminates and returns the following final state:

```
{pa : Logger |
  history: ({[n1 n0] [n2 n1 n0]} ... {[n1 n0] [n2 n0]})}
[n1 : router |...
  bestPath: (source: n1,dest: n0,pv:(n1 n0),metric: 2), ...]
[n2 : router |...
  bestPath: (source: n2,dest: n0,pv:(n2 n1 n0),metric: 1),...]
```

Note that this final state corresponds to one of the stable path assignment of Disagree described in Section 2, where node `n1` sets its best path to `[n1 n0]`, and node `n2` sets its best path to `[n2 n1 n0]`.

On the other hand, with the `rew` command which employs a different rewriting strategy, divergence scenario is simulated and route oscillation is observed in the simulation. This is because `frewrite` employs a depth-first position-fair rewriting strategy, while `rewrite` employs a left-most, outer-most strategy that coincides with the execution trace that leads to divergence.

5.2 Route Oscillation Detection

While Maude commands `frew/rew` explore a small portion of possible runs of the instance, the `search` command allows us to exhaustively explore the entire execution space. To exhaustively search BGP execution for route oscillation, we only need to first input the BGP instance’s network topology and routing policy to generate the corresponding initial state, as described in Section 4; and then use the `search` command to automatically search for oscillation. For example, for Disagree, we run:

```
search [1] gadget =>+ X such that detectCycle(X) = true .
```

Here, `gadget` is Disagree’s initial state, and `=>+ X` tells Maude to search for any reachable network state `X` such that at that state, the logger `pa` contains recurring path assignment (`detectCycle(X)=true`). `search` command exhaustively explores Disagree runs and returns with the first Disagree state that exhibits oscillation:

```
{pa : Logger | history: ({[n1 n2 n0] [n2 n0]}
                        {[n1 n2 n0] [n2 n1 n0]}
                        {[n1 n2 n0] [n2 n0]}
                        {[n1 n0] [n2 n0]})}
[n1 : router |...] [n2 : router |...] ...
```

Here, the resulting path assignment content in `pa` exhibits an oscillation (line 1, line 3).

In general, Maude allows us to exhaustively search for violation of a safety property `P` by running the following command:

```
search initialNetwork =>+ X:Configuration such that P(X) == false.
```

which tells Maude to exhaustively search for a network state X that violates P along all possible execution traces from the initial state `initialNetwork`. If Maude returns with `No solution`, we can conclude property P holds for all execution traces.

5.3 Case Studies

We summarize analysis that we have performed over well-known eBGP instances, including good gadget, bad gadget, disagree [10]. In addition, we analyze two iBGP configuration instances: a 9-node iBGP gadget [7] that is known to oscillate, and a 25-node configuration randomly extracted from the Rocketfuel [17] dataset. Rocketfuel is a well-known dataset on actual iBGP configurations that is made available to the networking community. Given that an ISP has complete knowledge of its internal router configurations, the Rocketfuel experiment presents a practical use case for using our tool to check an actual BGP configuration instance for safety.

For each BGP instance, we simulate its possible executions using rewriting commands (*Simulation*), and check for route oscillation using exhaustive search (*Exhaustive*). Our results are as follows:

	Disagree	Bad	Good	9-node iBGP	25-node iBGP
Simulation	2	NA	4	20	31
Exhaustive	2,10,No	181,641,No	10997,37692,Yes	20063,52264,No	723827,177483,Yes

Table 2: Summary of BGP analysis in Maude. In the first row, each entry shows the simulation time in milliseconds. In the second row, for each entry, the first value denotes exhaustive search time in milliseconds, the second denotes number of states explored, and the third on whether our tool determines the instances to be safe (“Yes”) or unsafe (“No”).

We have carried out these analysis on a laptop with 1.9 GB memory and 2.40GHz dual-cores running Debian 5.0.6. The version of Maude is Maude 2.4. While route oscillation detection explores the entire state space of the instance execution, the analysis time for rewriting based execution are measured for only one possible terminating execution (that converges to a stable path assignment).

Here are observations from our results: Single-trace simulation does not always terminate, and is helpful in finding permanent routing oscillation when Maude does not terminate (e.g., in executing Bad gadget ⁶), but can miss temporary oscillations which are only manifested on a particular executing trace. When Maude terminates, single-trace simulation time increases when network size grows. On the other hand, exhaustive search always provides a solid safety proof. For instances of similar network size, the search time for a safe instance (`good`) is considerably longer than that of an unsafe instance (`bad`). For instances of different sizes, as network size grows, exhaustive

⁶ Bad gadget always diverges and does not have any stable path assignment, therefore, when we simulate bad gadget with rewriting, Maude does not terminate, and we do not record the statistics.

search time grows exponentially. Nevertheless, even for the 25-node scenario, exhaustive search can be completed in 12 minutes. As future work, we are going to scale our analysis technique to larger networks.

6 Related Work

Maude is a widely used tool for a variety of protocol analysis. In addition to our use of Maude for analyzing BGP instances, there is also a huge literature of using Maude for other complex systems, such as security protocols [9], real-time systems [14], and active networking [5].

Theorem proving and model checking techniques have been applied to formal verification of network protocols. For instance, in [3], a routing protocol standard is formalized in the `Spin` model checker and `HOL` theorem prover, where `Spin` is used to check convergence of small network instances, and `HOL` is used to generalize the convergence proof for arbitrary network instances. Their work focuses on basic intra-AS routing protocols such as the distance-vector protocol, and does not consider policy interactions that occur in inter-AS routing protocols such as BGP. However, while our proofs are automated by Maude’s built-in simulation and exhaustive search capabilities, we are restricted to analyzing specific network instances. As future work, we plan to generalize our instance-based proofs towards more general properties on BGP stability, by leveraging Maude’s connection with existing theorem provers such as PVS [15].

Arye *et al.* [2] has attempted a similar formalization of eBGP gadgets in SPP using the Alloy [1] tool. Our approach differs from theirs in the overall goal of the formalization: Arye *et al.* uses Alloy to synthesize eBGP instances that exhibit certain behaviors such as divergence, whereas our approach takes an eBGP instance as input and analyzes it via simulation runs and exhaustive search. Our overall goal is to provide an easy-to-use library in Maude that eases the entire process of specifying and analyzing a BGP instance. Besides, in addition to eBGP gadgets, our library also supports iBGP instances and handles iBGP route ranking generation based on a separate IGP protocol 4.2.

7 Conclusion and Future Work

This paper presents our development of a Maude library for specifying and analyzing BGP instances. Our work aims to automate an important task for network designers when designing BGP protocols and safe policy guidelines. Our library uses Maude’s object-based specification language and enables the user to easily generate Maude specification by only requiring them to define the network topology and routing policies. To validate the feasibility of our library, we explored a variety of well-known BGP gadgets and an actual BGP instance obtained from the Rocketfuel dataset, and demonstrated the use of Maude’s analysis capabilities to detect possible divergence. All Maude code described in this paper is available at <http://netdb.cis.upenn.edu/discotec11>.

In addition to integrating our framework with the PVS theorem prover, our ongoing work includes: (1) more case studies on BGP instances and recent guidelines to explore the limits of our library, leading to possible extensions of our Maude library; and (2) releasing our tool for network designers to use.

8 Acknowledgment

The authors would like to thank Yiqing Ren and Wenchao Zhou for their valuable help in generating the iBGP instances (from the Rocketfuel dataset) used in our analysis. This research was supported by NSF grants CNS-0830949, CPS-0932397, IIS-0812270, TC-0905607, AFOSR Grant No: FA9550-08-1-0352, NSF CAREER CNS-0845552, and Office of Naval Research Grant N00014-10-1-0365.

References

1. Alloy. <http://alloy.mit.edu/community/>.
2. M. Arye, R. Harrison, and R. Wang. The Next 10,000 BGP Gadgets: Lightweight Modeling for the Stable Paths Problem. Princeton COS598D course project report.
3. K. Bhargavan, D. Obradovic, and C. A. Gunter. Formal verification of standards for distance vector routing protocols. *J. ACM*, 49(4):538–576, 2002.
4. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. *All About Maude: A High-Performance Logical Framework*. Springer, 2007.
5. G. Denker, J. Meseguer, and C. Talcott. Formal specification and analysis of active networks and communication protocols: The maude experience. *DARPA Information Survivability Conference and Exposition*, 2000.
6. N. Feamster, R. Johari, and H. Balakrishnan. Implications of Autonomy for the Expressiveness of Policy Routing. In *ACM SIGCOMM*, Philadelphia, PA, August 2005.
7. A. Flavel and M. Roughan. Stable and flexible iBGP. In *ACM SIGCOMM*, 2009.
8. L. Gao and J. Rexford. Stable internet routing without global coordination. In *SIGMETRICS*, 2000.
9. A. Goodloe, C. A. Gunter, and M.-O. Stehr. Formal prototyping in early stages of protocol design. In *Proc. ACM WITS'05*, 2005.
10. T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE Trans. on Networking*, 10:232–243, 2002.
11. C. Labovitz, G. Malan, and F. Jahanian. Internet Routing Instability. *TON*, 1998.
12. Maude. <http://maude.cs.uiuc.edu/>.
13. J. Meseguer. Conditional Rewriting Logic as a Unified Model of Concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
14. P. C. Ölveczky and J. Meseguer. Real-time maude: A tool for simulating and analyzing real-time and hybrid systems. *Electr. Notes Theor. Comput. Sci.*, 36, 2000.
15. PVS Specification and Verification System. <http://pvs.csl.sri.com/>.
16. M. Schapira, Y. Zhu, and J. Rexford. Putting BGP on the right path: A case for next-hop routing. In *HotNets*, Oct. 2010.
17. N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *SIGCOMM'02*.
18. L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica. HLP: A Next-generation Interdomain Routing Protocol. In *SIGCOMM*, 2005.
19. A. Wang, C. Talcott, L. Jia, B. T. Loo, and A. Scedrov. Analyzing BGP instances in maude. University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-11-08 2011. <http://netdb.cis.upenn.edu/papers/bgpMaudeTR.pdf>.