

A Simplified Formal Specification for BGP

Dan Pei¹, Dan Massey² and Lixia Zhang¹

UCLA¹
{peidan,lixia}@cs.ucla.edu

Colorado State University²
massey@cs.colostate.edu

UCLA CSD Technical Report TR040044

November 8, 2004

Abstract

As with those of other protocols, BGP protocol standard[1] specifies the protocol details using language description. However, language description can be ambiguous. In this technical report, we present a formal and short specification for BGP protocol such that most of the important BGP standard details are captured precisely using mathematic language. This report can be a quick and concise reference of BGP standard for those who have some basic BGP background.

In this report, we first define the several terminologies in Section 1. Then in Section 2, we define three categories of protocol functions at one router: operational environment, route filtering and aggregation policies, and path attribute manipulation policies. Then we use these functions to define the BGP protocol syntax and semantics requirements in Section 3.

1 Terminology Definitions

- AS Number: a 2-octet non-negative integer
IP address: a 4-octet non-negative integer
IP prefix: IP address/netmask, such as 131.179.0.0/16. Define $ipp_1 \leq ipp_2$ iff ipp_1 is included by ipp_2 .
- BGP path attribute set: attributes=(origin, as_path, next_hop, med, local_pref, atomic_aggregate, aggregator, community), such that:

origin : $\in \{IGP = 0, EGP = 1, INCOMPLETE = 2\}$
as_path : is a sequence($as_{seg_1}, as_{seg_2}, \dots, as_{seg_n}$), $n \geq 0$ such that:
 $\forall i, 1 \leq i \leq n$
 $as_{seg_i} = (AS_SET = 1, \{a_1, \dots, a_m\}), m \geq 1,$
 $a_j = \text{one AS Number}, 1 \leq j \leq m, \text{ where } \{a_1, \dots, a_m\} \text{ is an unordered set, OR}$
 $as_{seg_i} = (AS_SEQUENCE = 2, (b_1, \dots, b_l)), l \geq 1,$
 $b_j = \text{one AS Number}, 1 \leq j \leq l, \text{ where } (b_1, \dots, b_l) \text{ is an ordered sequence.}$
next_hop : IP address
med : four-octet non-negative integer
local_pref : four-octet non-negative integer
atomic_aggregate : $\in \{0, 1\}$
aggregator : AS Number:IP address
community : a set of four-octet values.

2 Functions Given by Administrators and Vendors

BGP provides vendors and administrators with a great flexibility. The vendor may select from a broad range of implementation options and the administrator can further refine these options and add additional local policies. For example, the BGP specification imposes some constraints on the route selection process, but leaves most decisions up to the vendors and administrators. One BGP router may choose to select routes based on the shortest AS path length, but a second BGP router may ignore the AS path length altogether and select routes based on the some other local policies. Both routers fully comply with the BGP specification. In order to capture the full range of BGP in mathematical description, the following functions must be provided by the vendors and/or administrators.

These given functions fall into three categories. The first category lists the AS numbers for the router and its peers, the corresponding IP addresses, and other local network topology information. The second category lists the route filtering and aggregation policies for the router. The third category lists policies for updating individual BGP path attributes.

2.1 Operational Environment

For a BGP router b , we are given the following:

- $Peers(b) = \{0, \dots, P\}$: a set of potential BGP peer routers. In our notation, Peer 0 will be used to represent static routes and routes learned from IGP.
- $as(x)$: the AS Number of x , $\forall x \in Peers(b) \cup \{b\}$. Define $as(0)=as(b)$.
- $interfaces(x)$: a subset of x 's interface IP addresses such that they are known by b . Define $interfaces(0) = interfaces(b)$.
- $ip(p)$: the IP address of p that is used in the peering session between b and p , where $ip(p) \in interfaces(p)$, $\forall p \in Peers(b) - \{0\}$.
- $ipb(p)$: the IP address of b that is used in the peering session between b and p , where $ipb(p) \in interfaces(b)$, $\forall p \in Peers(b) - \{0\}$.
- Function $same_subnet(addr_1, \dots, addr_n) \rightarrow \{0, 1\}$ such that $addr_1, \dots, addr_n$ are IP addresses and $same_subnet(addr_1, \dots, addr_2) = 1$ iff $addr_1, \dots, addr_n$ all share the same subnet.
- Function $multihop(b, p) = 0$ iff $\forall addr_1 \in interfaces(b), addr_2 \in interfaces(p), same_subnet(addr_1, addr_2) = 0$.
- $EBGP(p): \rightarrow \{0, 1\}$ such that $EBGP(p)=1$, iff $as(b) \neq as(p)$
- $Internalcost(addr)=$ the IGP cost from b to $addr$, which is an integer, and maximum value is $max_internal_cost$.
- $Length(aspath)=$ the length of $aspath$.
- $MinRouteAdver(p)$: the time between two successive advertisements(from b to p) of routes to a particular destination
- $MinRouteOrig(p)$: the time between two successive advertisements(from b to p) of routes to a destination in local AS

2.2 Route Filtering and Aggregation Policies

- Function $filter(p, direction, nlri, attributes) \rightarrow \{0, 1\}$ such that
 - $p \in Peers(b)$
 - $direction \in \{IN, OUT\}$
 - $nlri$ is an IP prefix
 - $attributes$ is a set of BGP attributes
 - $filter(p, IN, nlri, attributes) = 1$ iff the route $(nlri, attributes)$ will be accepted from peer p
 - $filter(p, OUT, nlri, attributes) = 1$ iff the route $(nlri, attributes)$ will be advertised to peer p
- $aggregates(b)$ = a set of IP prefixes to be aggregated by b
- Function $aggregate(agnlri, p, nlri_1, attributes_1, \dots, nlri_n, attributes_n) \rightarrow (attributes, include, exclude)$ such that
 - $agnlri \in aggregates(b)$
 - $p \in Peers(b)$
 - $\forall i, nlri_i \leq agnlri$, AND $nlri \in exclude$ or $nlri \in include$
 - $exclude \cap include = 0$
 - $\forall nlri_i \in include, attributes.med = attributes_i.med$
 - $\forall nlri_i \in include, attributes.nexthop = attributes_i.nexthop$
 - if $attributes.origin = IGP$, then $\forall nlri_i \in include, attributes_i.origin = IGP$
 - if $attributes.origin = EGP$, then $\exists nlri_i \in include$ such that $attributes_i.origin = EGP$ and $\forall nlri_i \in include, attributes_i.origin \in \{EGP, IGP\}$
 - if $attributes.origin = INCOMPLETE$, then $\exists nlri_i \in include$ such that $attributes_i.origin = INCOMPLETE$
 - $attributes.community$ = any set of community values
 - $attributes.aggregator = 0$, or $as(b) : addr$ such that $addr \in interfaces(b)$
 - $attributes.atomic_aggregate = 1$ if $\exists nlri_i \in includes$ such that $attributes_i.atomic_aggregate = 1$
 - $attributes.med$ =any med value
 - $attributes.local_pref$ =any local_pref value
 - $attributes.as_path = (asseq_1, \dots, asseq_n), n \geq 0$ such that
 - * if $\forall nlri_i, nlri_j \in include, attributes_i.as_path = attributes_j.as_path$, then $attributes.as_path = attributes_i.as_path, nlri_i \in include$
 - * if $a \in as_seq$ such that $(AS_SEQUENCE, as_seq) \in attributes.as_path$, then $\forall nlri_i \in include, a \in attributes_i.as_path$
 - * if $a \in as_set$ such that $(AS_SET, as_set) \in attributes.as_path$, then $\exists nlri_i \in include, a \in attributes_i.as_path$

- * if $a = a_i, 1 \leq i \leq m$, such that $(AS_SEQUENCE, (a_1, a_2, \dots, a_m)) = asseg_j$, then: $\forall b = a_k, i < k \leq m$ OR $b \in asseg_k, j < k \leq n$,
if $b = b_x, 1 \leq x \leq s$, such that $(AS_SEQUENCE, (b_1, b_2, \dots, b_s)) = asseg_t$ AND $1 \leq t \leq u, (asseg_1, asseg_2, \dots, asseg_u) = attributes_w.as_path, nlri_w \in include$, then $a = b_y, 1 \leq y < x$ OR $a \in asseg_z, 1 \leq z < t$.
- * if $a = a_i, 1 \leq i \leq m$, such that $(AS_SET$ OR $AS_SEQUENCE, (a_1, a_2, \dots, a_m)) = asseg_j$, then $a \neq a_k, \forall 1 \leq k \leq m, k \neq i$ AND $a \notin asseg_k, \forall 1 \leq k \leq n, k \neq j$.

2.3 Path Attribute Manipulation Policies

- In all of the following function definitions:
 - $p \in Peers(b), fp \in Peers(b)$
 - $nlri$ is an IP prefix
 - $direction \in \{IN, OUT\}$
 - $attributes$ is a set of BGP attributes
 - $function(0, OUT, nlri, attributes)$ is not defined.
- Function $setorigin(0, IN, nlri, attributes) \rightarrow \{0, 1, 2\}$ and $setorigin(p>0, IN/OUT, nlri, attributes)=attributes.origin$.
- Function $setmed(p, direction, nlri, attributes) \rightarrow$ non-negative integer
- Function $setlocalpref(p, direction, nlri, attributes) \rightarrow$ four-octet non-negative integer such that $setlocalpref(p, OUT, nlri, attributes)=\begin{cases} attributes.local_pref, if as(b) = as(p) \\ \emptyset, otherwise. \end{cases}$
- Function $setcommunity(p, direction, nlri, attributes) \rightarrow$ a set of four-octet values such that $setcommunity(p>0, IN, nlri, attributes)=attributes.community$
If $as(p) \neq as(b), attributes.med \neq 0$, then $setcommunity(p, OUT, nlri, attributes)$'s default value must be 0.
- Function $setaspath(p, direction, nlri, attributes) \rightarrow as_path$ such that
 - $setaspath(0, IN, nlri, attributes) = \emptyset, setaspath(p > 0, IN, nlri, attributes) = attributes.aspath$
 - if $as(p) = as(b)$ then $setaspath(p, OUT, nlri, attributes) = attributes.aspath$
 - if $as(p) \neq as(b)$ then $setaspath(p, OUT, nlri, attributes) = (as(b))^+ \cdot attributes.aspath$
- Function $setnexthop(p, fp, direction, nlri, attributes) \rightarrow$ IP address, where fp is the the peer from which the route is learned, such that
 - $setnexthop(0, 0, IN, nlri, attributes) \in interfaces(b), setnexthop(p > 0, fp = p, IN, nlri, attributes) = attributes.nexthop$.
 - if $as(p) = as(b)$ then $setnexthop(p, OUT, nlri, attributes) = attributes.nexthop$ OR $setnexthop(p, OUT, nlri, attributes) \in interfaces(b)$.
 - if $as(p) \neq as(b) AND multihop(p, b) = 0$ then $same_subnet(setnexthop(p, fp, OUT, nlri, attributes), ip(p)) = 1$ and one of:

- * $setnexthop(p, fp, OUT, nlri, attributes) = ipb(p)$ OR
 - $subset(ip(p), addr) = 1$ AND $addr \in interfaces(b)$ AND $setnexthop(p, fp, OUT, nlri, attributes) = addr$
 - $as(fp) = as(b)$ AND $setnexthop(p, fp, OUT, nlri, attributes) \in interfaces(fp)$
 - $as(fp) \neq as(b)$ AND $setnexthop(p, fp, OUT, nlri, attributes) = attributes.nexthop$ and $\exists addr \in interfaces(b)$ such that $same_subnet(addr, ip(p)) = 1$
- if $as \neq as(b)$, $multihop(p, b) = 1$ then either:
 - * $setnexthop(p, fp, OUT, nlri, attributes) = attributes.nexthop$
 - * $setnexthop(p, fp, OUT, nlri, attributes) = ipb(p)$.

3 Protocol Definition

This section will describe the protocol semantic definition and we assume that several protocol syntactic functions will take care of the tasks such as maintaining peering section, reporting peer status, checking the format of incoming update packets, forming the outbound update packets.

3.1 Protocol Syntactic Functions and Environment

- T=the current time, and it is a integer.
- Peering Session Function:

$$UP(p > 0, T) = \begin{cases} 1, \text{ peering session between } b \text{ and } p \text{ is in ESTABLISHED STATE at time } T, \\ 0, \text{ otherwise.} \end{cases}$$

Define $UP(0, T) = 1$. $Up(p, T)$ also maintains the peering session, including sending /processing OPEN/KEEPALIVE/NOTIFICATION messages.

- BGP route update= $(wnlri, anlri, attributes, time)$, such that
 - $wnlri$: a set of IP prefixes of the routes to be withdrawn
 - $anlri$: a set of IP prefixes of the routes to be announced ($wnlri \cap anlri = \emptyset$)
 - $attributes$: BGP path attribute set.
 - $time$: the time the route is received or sent
- Any changes of the routes learned from Peer 0 (static routes or routes learned from IGP) will be translated to a BGP update format, applying any applicable given functions.
- Last_Recv(p,T)=the last UPDATE message that b received from p prior to time T
- Time T will increase by 1 after any of the following:
 - Given Functions changed(including policy changes)
 - receiving an update from a peer
 - $UP(p, T)$ function changes

3.2 Protocol Semantic Definition

These functions are conceptual, the implementation may choose any algorithm as long as they conform to these definitions.

- $\text{Import}(p, \text{nlri}, \text{attributes}) \rightarrow$ path attributes such that $\text{attributes} =$

$$\left\{ \begin{array}{ll} \emptyset, & \text{if } \text{as}(b) \in \text{attributes.as_path}; \\ \emptyset, & \text{if } \text{filter}(p, \text{IN}, \text{nlri}, \text{attributes}) = 1; \\ \{\text{setorigin}(p, \text{IN}, \text{nlri}, \text{attributes}), \\ \text{setaspath}(p, \text{IN}, \text{nlri}, \text{attributes}), \\ \text{setnexthop}(p, p, \text{IN}, \text{nlri}, \text{attributes}) \\ \text{setmed}(p, \text{IN}, \text{nlri}, \text{attributes}), \\ \text{setlocalpref}(p, \text{IN}, \text{nlri}, \text{attributes}), \\ \text{attributes.atomic_aggregate}, \\ \text{attributes.aggregator}, \\ \text{attributes.community}\}, & \text{otherwise} \end{array} \right.$$

- $\text{Select}(p, \text{attributes}): \rightarrow$ real number such that:

- $\text{Select}(p, \text{attributes}) = 0$, iff $\text{RIB-Local}(\text{attributes.nexthop}, \text{T}) = \emptyset$
- $\text{Select}(p, \text{attributes}) = \text{attributes.local_pref} + (1 - \frac{\text{Length}(\text{attributes.aspath})}{\text{max_aspath_length} + 1}) + (1 - \frac{\text{attributes.origin}}{\text{max_origin} + 1}) * 10^{-1} + (1 - \frac{\text{attributes.med}}{(\text{max_med} + 1)}) * 10^{-2} + \frac{\text{EBGP}(p)}{2} * 10^{-3} + (1 - \frac{\text{Internalcost}(\text{attributes.nexthop})}{(\text{max_internal_cost} + 1)}) * 10^{-4} + (1 - \frac{\text{router_id}}{(\text{max_router_id} + 1)}) * 10^{-5}$

- Export Policy function: $\{p, \text{fp}, \text{nlri}, \text{attributes}\} \rightarrow$ path attributes such that $\text{Export}(p, \text{fp}, \text{nlri}, \text{attributes}) =$

$$\left\{ \begin{array}{ll} \emptyset, & \text{if } \text{ip}(p) = \text{attributes.next_hop}; \\ \emptyset, & \text{if } \text{filter}(p, \text{OUT}, \text{nlri}, \text{attributes}) = 1; \\ \{\text{setorigin}(p, \text{OUT}, \text{nlri}, \text{attributes}), \\ \text{setaspath}(p, \text{OUT}, \text{nlri}, \text{attributes}), \\ \text{setnexthop}(p, \text{fp}, \text{OUT}, \text{nlri}, \text{attributes}), \\ \text{setmed}(p, \text{OUT}, \text{nlri}, \text{attributes}), \\ \text{setlocalpref}(p, \text{OUT}, \text{nlri}, \text{attributes}), \\ \text{attributes.atomic_aggregate}, \\ \text{attributes.aggregator}, \\ \text{setcommunity}(p, \text{OUT}, \text{nlri}, \text{attributes})\}, & \text{otherwise.} \end{array} \right.$$

- $\text{Adj-RIB-In}(p, \text{nlri}, \text{T}) \rightarrow$ path attributes such that $\text{attributes} =$

$$\left\{ \begin{array}{ll} \emptyset, & \text{if } \text{UP}(p, \text{T}) = 0; \\ \emptyset, & \text{if } \text{UP}(p, \text{T}) = 1 \text{ AND } \text{nlri} \in \text{Last_Recv}(p, \text{T}).\text{wnlri}; \\ \text{Last_Recv}(p, \text{T}).\text{attributes}, & \text{if } \text{UP}(p, \text{T}) = 1 \text{ AND } \text{nlri} \in \text{Last_Recv}(p, \text{T}).\text{anlri}, \\ \text{AdjRIBIn}(\text{nlri}, p, \text{T} - 1), & \text{otherwise.} \end{array} \right.$$

- $\text{RIB-Local}(nlri, T) = (fp, \text{attributes}) = (p_{best}, \text{Adj-RIB-In}(p_{best}, nlri, T))$, $p_{best} \in \text{peers}(b)$ such that:
 $\text{select}(p_{best}, \text{Adj-RIB-In}(p_{best}, nlri, T)) = \max_{p \in \text{peers}(b)} \{\text{select}(p, \text{Import}(p, nlri, \text{Adj-RIB-In}(p, nlri, T)))\}$
 > 0 .
- $\text{Adj-RIB-Out}(p, nlri, T) \rightarrow$ path attributes such that
 - $\text{Adj-RIB-Out}(p, nlri, T) = \text{Export}(p, \text{RIB-Local}(nlri, T).fp, \text{RIB-Local}(nlri, T).attributes)$ iff $nlri \notin \text{aggregates}(b)$ OR
 - $\text{Adj-RIB-Out}(p, nlri, T) = \text{agattributes}$, if $nlri \in \text{aggregates}(b)$ such that:
 - * $\text{aggregate}(nlri, p, nlri_1, \text{attributes}_1, \dots, nlri_n, \text{attributes}_n) = (\text{agattributes}, \text{include}, \text{exclude})$,
and
 - * $|\text{include}| \geq 2$.
- $\text{setatomicaggregate}(p, nlri, \text{attributes}) =$

$$\begin{cases} 1, & \text{if } \text{attributes.atomic_aggregate} = 1 \\ 1, & \text{if } nlri \in \text{aggregates}(b), \text{ and } \exists nlri_{in} \leq nlri, \text{ such that :} \\ & \text{Adj-RIB-In}(p, nlri_{in}, T) \neq \emptyset, \text{ AND } \text{Adj-RIB-Out}(p, nlri_{in}, T) = \emptyset \\ 0, & \text{otherwise} \end{cases}$$
- Adj-RIB-Out change sequence $\text{RibOutSeq}(p, nlri, T)$ is a sequence of path attributes:
 $(\text{attributes}_1, \dots, \text{attributes}_m)$ such that:
 - $\forall 1 \leq i \leq m, \text{attributes}_i = \text{Adj-RIB-Out}(p, nlri, t_i), t_1 < \dots < t_m \leq T$, AND
 - $\forall 1 \leq i < m, \text{attributes}_i \neq \text{attributes}_{i+1}$ AND
 - $\forall 1 \leq i < m$, there is no $\text{Adj-RIB-Out}(p, nlri)$ changes between t_i and t_{i+1} or between t_m and T .
- Advertisement sequence $\text{AdvSeq}(p, nlri, T) = (a_1, \dots, a_n)$, such that:
 - a_i is the update that b sent to p , AND
 - $\forall i$, EITHER $nlri \in a_i.anlri$, OR $nlri \in a_i.wnlri$, AND
 - $a_1.time < \dots < a_n.time \leq T$, AND
 - $\forall 1 \leq i < n$, there is no update a from b to p such that $nlri \in a.nlri$ AND:
 $a_i.time < a.time < a_{i+1}.time$ OR $a_n.time < a.time < T$

3.3 Requirements

$\forall p > 0, nlri, T, \text{AdvSeq}(p, nlri, T) = (a_1, \dots, a_n)$ must satisfy:

R 1 Advertisement Frequency Requirements:

- if $as(b) \neq as(p)$, then: $\forall 1 \leq i < n$, if $nlri \in a_i.anlri$ AND $nlri \in a_{i+1}.anlri$, then
 $a_{i+1}.time - a_i.time \geq \text{MinRouteAdver}(p) * 0.75$.

- $\forall 1 \leq i < n$, if $a_{i+1}.attributes.aspath = \emptyset$, then:
 $a_{i+1}.time - a_i.time \geq MinRouteOrig(p) * 0.75$.

R 2 $\exists B$, such that:

- if $\forall i \leq B$, $Adj-RIB-Out(p, nlri, T) = Adj-RIB-Out(p, nlri, C-i) \neq \emptyset$, then $nlri \in a_n.anlri$, AND $a_n.attributes$
 $Adj-RIB-Out(p, nlri, T).attributes$.
- if $\forall i \leq B$, $Adj-RIB-Out(p, nlri, T) = Adj-RIB-Out(nlri, C-i) = \emptyset$, then $nlri \in a_n.wnlri$, AND
 $nlri \in a_{n-1}.nlri$.

R 3 Sequence $(a_1.attributes, \dots, a_n.attributes)$ must be a subsequence of $RibOutSeq(p, nlri, T)$.

References

- [1] Y. Rekhter, T. Li, and S. Hares. Border Gateway Protocol 4. <http://www.ietf.org/internet-drafts/draft-ietf-idr-bgp4-26.txt>, Oct 2004.