Wireless Networks 10, 595–605, 2004

BGP-S: A Protocol for Terrestrial and Satellite Network Integration in Network Layer

EYLEM EKICI *
Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210, USA

CHAO CHEN
Broadband and Wireless Networking Laboratory, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

Abstract. To accomplish network layer integration of terrestrial and satellite IP networks, special exterior gateway protocols are needed. In this work, a new exterior gateway protocol called *Border Gateway Protocol – Satellite version* (BGP-S) is introduced that enables automated discovery of paths that go through the satellite network. This protocol is designed to work in only one terrestrial gateway in every Autonomous System and enables the forwarding of discovered paths in the Internet using the BGP-4 protocol. The performance of BGP-S is investigated through simulations.

Keywords: IP-based routing, satellite networks, BGP-4, exterior gateway protocol

1. Introduction

Satellite networks are becoming increasingly important for global communications. With the explosive growth of the Internet, the IP technology is being pushed to the satellite networks. To realize this, satellites carry IP-switches that forward packets independently. These IP-switches are connected to each other as well as to ground stations. Several issues related to IP-based satellite networks have been reviewed in [1]. Routing in the LEO satellite environment is a challenging problem because of the dynamic nature of the satellite networks. In recent years, several routing algorithms and protocols have been proposed for IP-based LEO satellite networks [2–5].

However, the use of the IP-based satellite networks as a part of the Internet cannot be accomplished only by solving the routing problem of the satellite networks. The integration of the IP-based satellite networks must assure their interoperability with the terrestrial IP networks. Previously, satellite network integration issues have been pointed out in [1,6,7]. As suggested in these papers, the satellite network can be viewed as a separate Autonomous System (AS) with a different addressing scheme. To reduce the load on the satellite network, terrestrial gateways act as border gateways on behalf of the satellite network and perform the address translations. Then, paths over both networks can be discovered using an Exterior Gateway Protocol such as BGP [8]. However, since the internal and external metrics for terrestrial ASs and the satellite network are different, special care must be taken. None of the studies mentioned above provides a detailed solution as how this network level integration can be accomplished.

In this work, we propose the *Border Gateway Protocol – Satellite version* (BGP-S). The satellite network is considered as an AS with special properties. BGP-S is designed to coexist with the BGP-4 [8] and support automated discovery of paths that include the satellite hops. It is designed to be implemented in only one terrestrial gateway in every AS that is connected to the satellite network. Since the delay in the satellite network can be much longer than in a terrestrial AS, the acceptance of paths involving satellite hops is accomplished through active delay measurements. The general hybrid network architecture is introduced in section 2. In section 3, details of packet forwarding scheme is explained. The BGP-S protocol is described in detail in section 4. In section 5, an integrated terrestrial/satellite network model is built, and the performance of BGP-S is evaluated on this model. Finally, section 6 concludes this paper.

2. The hybrid terrestrial/satellite network architecture

The general hybrid network consists of the terrestrial Internet and an IP-based satellite network. The terrestrial Internet is organized into Autonomous Systems (ASs). Inside every AS, the routing is accomplished through *Interior Gateway Protocols* (IGPs). The inter-AS routing is based on an *Exterior Gateway Protocol* (EGP), specifically, *Border Gateway Protocol version 4* (BGP-4) [8]. The satellite network should carry the following properties:

- The satellite network should be able to forward individual data packets between two gateways on the Earth. The satellite network may use its own native packet formats and its own addressing scheme.

* Corresponding author. Eylem Ekici was with the Broadband & Wireless Networking Laboratory, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 when this work was performed.
E-mail: ekici@ece.osu.edu

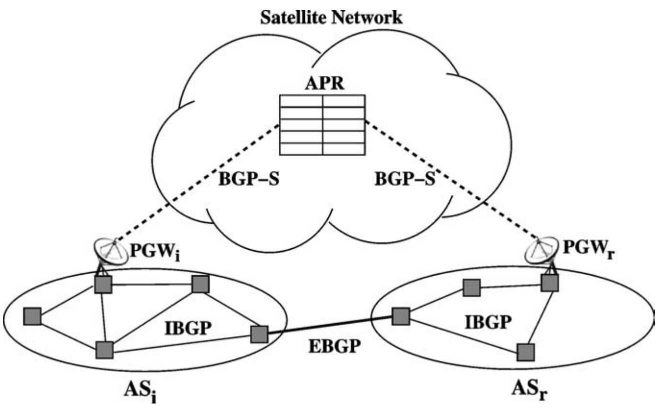

Figure 1. The hybrid terrestrial/satellite network architecture.

- There is no constraint on the satellite topology as long as any two terrestrial gateways can be connected over the satellite network. The satellite network can consist of any number of satellites in one or more orbits as long as every terrestrial gateway is always in the coverage area of at least one satellite and there exist a path to every other terrestrial gateway.
- There is no constraint on the routing protocol used in the satellite network, i.e., any custom routing protocol with static or dynamic routing tables/strategies is acceptable.

A sample structure of the hybrid terrestrial/satellite network is shown in figure 1. In this figure, two autonomous systems, AS_i and AS_r, are depicted. The autonomous systems are connected to the satellite network via a gateway. AS_i and AS_r are also connected with terrestrial links. Note that this figure is only a partial view of a likely network topology. There may be more autonomous systems with possibly different number of gateways and connected in a more complex way.

The following is a list of notations used in this work:

Autonomous system. The collection of routers under the same technical and administrative control is referred to as an Autonomous System. The autonomous systems are denoted by AS_i as shown in figure 1.

Routers and BGP speakers. The routers in every autonomous system AS_i are denoted by $R_{i,j}$, for $j = 0, \ldots, \mathbf{N}_i^R - 1$, where $\mathbf{N}_i^R$ is the number of routers in AS_i. The BGP speakers are the routers that implement BGP and they are denoted by $BS_{i,j}$, for $j = 0, \ldots, \mathbf{N}_i^{BS} - 1$, where $\mathbf{N}_i^{BS}$ is the number of BGP speakers in AS_i. Note that $\mathbf{N}_i^{BS} \leqslant \mathbf{N}_i^R$ and $\{BS_{i,j}\} \subseteq \{R_{i,j}\}$.

Network address. A network address NA_i is the longest common IP prefix shared by the network elements in that subnetwork. An example network address is `193.140.196.0/24`.

AS path. An AS path $\mathbf{P}_{AS_i}^j(NA_k)$ is an ordered list of autonomous systems $(AS_i, \ldots, AS_x)$, which is the jth alternative path for AS_i to reach the network address NA_k, where NA_k resides in AS_x.

Gateways and peer gateways. The gateways are the terrestrial stations that enable the communication between the autonomous systems and the satellite network. In an autonomous system AS_i, the number of gateways is $\mathbf{N}_i^{GW}$, and the gateways are denoted by $GW_{i,j}$, for $j = 0, \ldots, \mathbf{N}_i^{GW} - 1$. One of the gateways is designated as the peer gateway and implements the BGP-S protocol used for path discovery over the satellite network. The peer gateway in an autonomous system AS_i is denoted by PGW_i as shown in figure 1. A peer gateway is a gateway, a router, and a BGP speaker at the same time.

Active peer register. The active peer register (APR) is the list of active peer gateways connected to the satellite network. APR can be maintained on the Earth as well as in the satellite network, where it can be reached by peer gateways over pre-configured paths. APR can also be duplicated as long as all copies are updated in real-time.

In addition to these components, there are also other components in the hybrid network. The terrestrial network contains routers and hosts, and there are satellites with on-board routers. Satellites are denoted by S_i, for $i = 0, \ldots, \mathbf{N}^S - 1$, where $\mathbf{N}^S$ is the number of satellites. Note that no specific satellite constellation or organization of the satellites is assumed. Thus, only the index i in S_i is sufficient to refer to a specific satellite.

3. Packet forwarding

The packet forwarding from one terrestrial gateway to the next occurs with "IP over IP" tunneling in the satellite network. Under this scheme, packets are encapsulated individually into native satellite packets before they are sent to the satellite network by the terrestrial gateway. Native satellite packets carry the address of the next terrestrial gateway which can be interpreted by all satellites in the network. Hence, satellites do not need to keep track of all IP addresses. The satellite network is responsible for relaying the packets between terrestrial gateways only. It is assumed that the addressing scheme used by the satellite network and the mapping of these addresses to IP-addresses are available in the terrestrial gateways.

The packet processing in the terrestrial gateways is the most important step to use the satellite network as a part of IP paths. While the routers in the terrestrial network continue using the standard packet forwarding procedures, the terrestrial gateways must translate the IP addresses and encapsulate the IP packets into native satellite packets. For this procedure to work, the terrestrial gateways must be addressable both by the terrestrial and satellite network.

Definition 1 (Next hop function NH). Let P denote a packet received by a terrestrial gateway. The function $NH(P)$ returns the next hop on the path of the packet P towards its destination.

Definition 2 (Satellite next hop function SNH). Let P denote a packet received by a terrestrial gateway $GW_{i,j}$, and the next hop for packet P be a terrestrial gateway $GW_{r,s}$, i.e., $NH(P) = GW_{r,s}$, where $(r, s) \neq (i, j)$, which is reachable through satellite network. The function $SNH(GW_{r,s})$ returns the satellite S_t, where $GW_{i,j}$ should first send the packet P to such that P reaches $GW_{r,s}$.

Upon receiving a packet P, a terrestrial gateway $GW_{i,j}$ processes the packet as follows:

1. The gateway determines the next hop $NH(P)$ for the received packet P.
2. If the packet's next hop is not a terrestrial gateway, i.e., $NH(P) \notin \{GW_{r,s} \mid (r, s) \neq (i, j)\}$, it forwards the packet to the next hop without any modifications.
3. If the next hop of the packet P is a terrestrial gateway, i.e., $NH(P) = GW_{r,s}$, $(r, s) \neq (i, j)$, then P is encapsulated into a native satellite packet with $GW_{r,s}$ as the destination and sent to its next hop S_t in the satellite network, where $S_t = SNH(GW_{r,s})$.

Note that it is assumed that no two terrestrial gateways are connected to each other with terrestrial links. If it is the case, then the function *NH* should be modified such that it also indicates if the next hop should be reached through the satellite network or over a direct terrestrial link. When a terrestrial gateway receives a native satellite packet from a satellite, it simply extracts the payload from the satellite packet and processes it as a regular IP packet.

4. The BGP-S protocol

To allow the automated discovery of paths that pass through the satellite network, we introduce a new protocol called the *Border Gateway Protocol – Satellite version* (BGP-S). BGP-S has the same basic functionality as the BGP-4 [8], which means that the AS policies used in BGP-4 are adopted to control routing traffic among networks. However, using the BGP-S together with BGP-4 has two main advantages. First of all, the satellite network does not directly participate in the path calculations. Instead, it is only responsible for carrying data packets and (possibly) keeping track of the active peer gateways. Hence, the complexity added to the satellite system is kept at a minimum. Secondly, if the satellite network is regarded as a regular autonomous system, there would not be any difference between a terrestrial AS and the satellite network. This may be misleading in many cases since the delays in the satellite network are much larger than in a terrestrial AS. Therefore, under BGP-4, if one or more satellite hops are involved in the AS-path, it is necessary to manually configure the routing strategies according to the location of the ASs and delay estimations. The BGP-S eliminates the need for manual configuration and enables automatic adaptation based on the delays in the satellite and terrestrial networks.

In the hybrid network model, BGP-4 and BGP-S are used together as shown in figure 1, where APR is located in the satellite network. Between the terrestrial BGP speakers, the BGP-4 protocol is used. More specifically, the Interior-BGP (IBGP) is used among the BGP speakers in the same AS. The BGP speakers that belong to different ASs use Exterior-BGP (EBGP). Although the message formats are the same both for IBGP and EBGP, there are differences in message processing. Peer gateways communicate over the satellite network using the BGP-S protocol. Peer gateways must implement both BGP-4 and BGP-S.

There are two important rules in a system implementing the BGP-S:

Rule 1. There is only one peer gateway in an AS.
Rule 2. The routing policies that are configured for the BGP-4 are automatically adopted by BGP-S.

The first rule aims to limit the number of peer gateways to the number of ASs directly connected to the satellite network. Furthermore, it eliminates duplication of information received in an AS. The second rule ensures that BGP-S is fully compatible with the BGP-4 protocol, hence with the existing Internet infrastructure. These rules may include elimination of paths that contain certain ASs, ensuring that transit traffic is not carried, etc. Detail description of BGP-4 can be found in [8,9]. The details of the BGP-S protocol are provided in the following sections.

4.1. BGP-S connection setup

The BGP-S protocol uses TCP connections between two peer gateways for communication. A BGP-S connection is closed either by an explicit *NOTIFICATION* message or when no messages are received from the other party within a predetermined time-out period. Considering the number of active peer gateways, the time-out period is suggested to be longer than in BGP-4, approximately 10 seconds. The connection setup is accomplished through the following steps, as also shown in figure 2:

1. When a peer gateway PGW_i becomes active and wants to connect to other peer gateways, it sends an $Alive(PGW_i)$ message to the Active Peer Register (APR).
2. The APR sends a list of already active peer gateways to PGW_i.
3. PGW_i acknowledges the reception of the active peer gateway list to the APR.

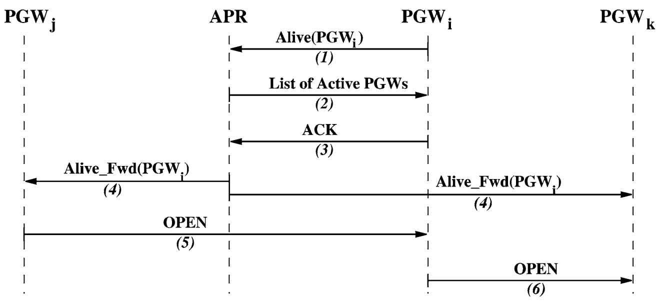

Figure 2. The activation of peer gateways and connection setup.

4. The APR sends to all other active peer gateways the *Alive_Fwd*(PGW_i) to notify them about the availability of the peer gateway PGW_i.
5. If an already active peer gateway PGW_j wants to establish a BGP-S connection, then it sends an *OPEN* message to PGW_i.
6. PGW_i can establish a BGP-S connection to any other peer gateway PGW_k in the active peer register by sending an *OPEN* message.

The *Alive*(PGW_i) message contains the IP and satellite network addresses of the peer gateway PGW_i as well as the AS number where PGW_i resides. The *Alive_Fwd*(PGW_i) message contains the same information as the *Alive*(PGW_i) message. The difference is that *Alive*() messages are created by the peer gateways that become active, and *Alive_Fwd*() messages are created by the APR to notify other peer gateways of the availability of a new peer gateway. The *OPEN* message has the same format as in the BGP-4 protocol.

4.2. Path discovery and prioritization

A peer gateway learns paths both via BGP-S and BGP-4. If it decides to advertise the paths to other peer gateways over BGP-S, then it uses *UPDATE* messages which have the same format as in BGP-4. It is important to note that the paths learned via BGP-S cannot be processed like the paths learned through BGP-4. The reasons for this differentiation were presented at the beginning of section 4. While processing these paths, it is important to be consistent with policies configured with the BGP-4 protocol. Then, the paths are compared based on the delay to the target network. Note that the delay comparison is just an approximation of the real-time delay. The delay changes continuously due to fluctuations in the traffic load and it is not feasible to check the delay to all possible network addresses periodically. In order to discover the delay to a given network, the following new messages are used:

***POLL*() message.** The *POLL* message is used to request a delay measurement to a specified network or network element. *POLL*(PGW_i, PGW_j, A) is a message sent by the peer gateway PGW_i to PGW_j to learn about the delay between PGW_j and A, where A can be a network or a network element. Every *POLL* message contains the message creation timestamp.

***DELAY*() message.** The *DELAY*() message is a reply to a *POLL*() message. The *DELAY*(PGW_j, PGW_i, A, B, d) is a message sent by the peer gateway PGW_j to PGW_i telling that the delay between itself and an network element B in the network A is d. If A is a network element, then $A = B$.

If $A = PGW_j$, then *DELAY*() is like a ping response; the receiving peer gateway PGW_j replies immediately with a delay equal to the timestamp in the *POLL*() message. Then, the peer gateway PGW_i calculates the round trip delay to PGW_j. If A is a network address, then PGW_j measures the delay to the network element B in the network A. Then the *DELAY*() message contains this measured delay as d. When the delay to a network A is needed, PGW_j selects a network element B in the network A and measures the delay from itself to B. The delay can be measured using the ping utility. However, any other method can be used for delay measurement, as well.

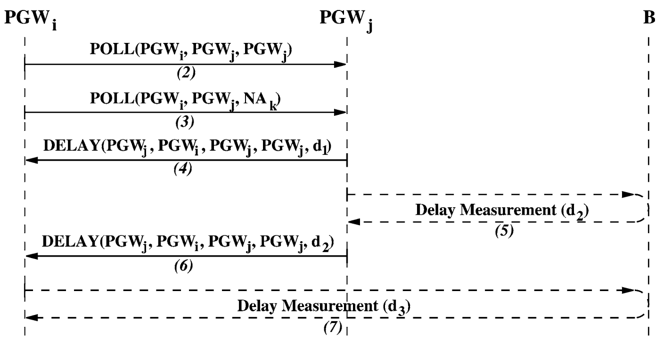

Figure 3. The processing of AS paths learned via BGP-S.

4.2.1. New path discovery via BGP-S

Assume a peer gateway PGW_i learns from PGW_j via BGP-S the AS path $\mathbf{P}_{AS_j}(NA_k)$ to reach the network NA_k. The new AS path $\mathbf{P}_{AS_j}(NA_k)$ is processed following the steps below, which are also shown in figure 3.

1. PGW_i checks $\mathbf{P}_{AS_j}(NA_k)$ with the policies setup for BGP-4 protocol. If there is a conflict, then $\mathbf{P}_{AS_j}(NA_k)$ is discarded.
2. If $\mathbf{P}_{AS_j}(NA_k)$ conforms with the BGP-4 policies and the delay from PGW_i to PGW_j is not available to PGW_i, then PGW_i sends a *POLL*(PGW_i, PGW_j, PGW_i) message to PGW_j.
3. PGW_i also sends a *POLL*(PGW_i, PGW_j, NA_k) message to PGW_j to learn the delay between PGW_j and the network NA_k.
4. PGW_i receives the *DELAY*(PGW_j, PGW_i, PGW_j, PGW_j, d_1) message from PGW_j. The delay d_1 to PGW_j, is estimated as the half of the difference of the current time T_{cur} and the timestamp d, i.e., $d_1 = (T_{cur} - d)/2$.
5. PGW_j measures the delay d_2 to the network element B in network NA_k.
6. PGW_i receives the *DELAY*(PGW_j), PGW_i, NA_k, B, d_2) message from PGW_j.
7. PGW_i measures the delay d_3 to B if there exists an AS path $\mathbf{P}_{AS_i}(NA_k)$ to reach the network NA_k in the Routing Information Base (RIB) of BGP-4. If there is no such entry in the BGP-4 RIB, then the delay to B is assigned infinity, i.e., $d_3 = \infty$.
8. If d_3 is infinity, then $\mathbf{P}^1_{AS_i}(NA_k)$ is created by appending AS_i to $\mathbf{P}_{AS_j}(NA_k)$ and inserted to BGP-4 RIB with a default local preference value.
9. Assume that there is already an AS path $\mathbf{P}^*_{AS_i}(NA_k)$ used in AS_i to reach the network NA_k such that $\mathbf{P}^*_{AS_i}(NA_k) =$

$\arg\max_{X \in \{\mathbf{P}_{AS_i}(NA_k)\}} LocalPref(X)$, where the function $LocalPref(X)$ gives the local preference value of the AS path X. If $d_1 + d_2 \geqslant d_3$, i.e., the new path over the satellite network is longer than the already available AS path, then the new AS path $\mathbf{P}^{p+1}_{AS_i}(NA_k)$ is inserted into BGP-4 RIB with a local preference value of $LocalPref(\mathbf{P}^{*}_{AS_i}(NA_k)) - 1$, where p is the number of AS paths to NA_k already in the RIB.

10. Under the same conditions as in the previous step, if $d_1 + d_2 < d_3$, i.e., the new path over the satellite network is shorter, then the new AS path $\mathbf{P}^{p+1}_{AS_i}(NA_k)$ is inserted into BGP-4 RIB with the local preference value of $LocalPref(\mathbf{P}^{*}_{AS_i}(NA_k)) + 1$, where p is the number of AS paths to NA_k already in the RIB.

When an AS path is inserted into the BGP-4 RIB by a peer gateway, the delay information remains local to the BGP-S protocol. The delay comparison is advertised to the BGP speakers in the same network implicitly with the local preference value, which is propagated with the new path information. Note that the local preference values of the AS paths inserted by BGP-S are related with the existing AS paths in the RIB. Although a relative local preference assignment is not allowed under BGP-4, BGP-S assigning relative local preference values does not affect the integrity of the BGP-4 because there is only one network entity per AS that is allowed to perform this operation.

4.2.2. New path discovery via BGP-4

Assume that a new AS path $\mathbf{P}^{p+2}_{AS_i}(NA_k)$ is advertised via BGP-4, which has a higher local preference value than the currently used, i.e.,

$$LocalPref\big(\mathbf{P}^{p+2}_{AS_i}(NA_k)\big) > LocalPref\big(\mathbf{P}^{*}_{AS_i}(NA_k)\big).$$

Also let $\mathbf{P}^{q}_{AS_i}(NA_k)$ be the AS path with the best delay performance to the network NA_k among the AS paths learned via BGP-S. The peer gateway PGW_i performs the following steps to process the new AS path:

1. If the BGP-4 RIB does not contain any path to NA_k that was learned via BGP-S, then no action is taken.
2. Otherwise, the delay to the network element B in NA_k is measured for $\mathbf{P}^{p+2}_{AS_i}(NA_k)$ and $\mathbf{P}^{q}_{AS_i}(NA_k)$. The measurements are taken following the steps 2–7 in section 4.2.1, obtaining the delays d_1, d_2, and d_3.
3. If $d_1 + d_2 \geqslant d_3$, i.e., the AS path over the satellite network $\mathbf{P}^{q}_{AS_i}(NA_k)$ is longer than the new AS path $\mathbf{P}^{p+2}_{AS_i}(NA_k)$, then no action is taken.
1. If $d_1 + d_2 < d_3$, i.e., the AS path over the satellite network $\mathbf{P}^{q}_{AS_i}(NA_k)$ is shorter than the new AS path $\mathbf{P}^{p+2}_{AS_i}(NA_k)$, then PGW_i updates the local preference of $\mathbf{P}^{q}_{AS_i}(NA_k)$ as $LocalPref(\mathbf{P}^{p+2}_{AS_i}(NA_k)) + 1$. Then, PGW_i advertises the path $\mathbf{P}^{q}_{AS_i}(NA_k)$ with the updated local preference value.

Note that the delay of the paths over the satellite network is re-measured by the peer gateways when learning new paths over BGP-4. However, regular delay monitoring of all the paths by peer gateways would not be feasible given the number of the ASs we consider and the high protocol overhead it would introduce.

4.2.3. Path withdrawal

When a path is withdrawn either via BGP-4 or BGP-S, the peer gateway PGW_i in AS_i must check the RIB and possibly modify the local preference value of the shortest AS path that goes over the satellite network. Assume that there are p paths in the BGP RIB to reach the network NA_k. Upon receiving an *UPDATE* message that contains the withdrawal of an AS path that leads to NA_k, the peer gateway PGW_i performs the following operations:

1. If the withdrawn AS path is not the one that is currently used, no action is taken.
2. If the currently used path is withdrawn and the AS path with the next highest local preference value is learned via BGP-4, then no action is taken.
3. If the AS path with the next highest local preference value is learned over BGP-S, then the AS path that is learned via BGP-4 and has the largest local preference value is found, which we call $\mathbf{P}^{t}_{AS_i}(NA_k)$.
4. All AS paths with larger local preference values than $\mathbf{P}^{t}_{AS_i}(NA_k)$ are collected in the set $\mathbf{P}^{Sat}_{AS_i}(NA_k)$.
5. The delays of all AS paths in $\mathbf{P}^{Sat}_{AS_i}(NA_k)$ are measured as described in section 4.2.1, steps 2–7. The delay of $\mathbf{P}^{t}_{AS_i}(NA_k)$ is also measured as described in these steps.
6. Let us assume that the AS path $\mathbf{P}^{s}_{AS_i}(NA_k)$ has the lowest delay d_s among all paths in $\mathbf{P}^{Sat}_{AS_i}(NA_k)$. Also assume that the delay of $\mathbf{P}^{t}_{AS_i}(NA_k)$ is d_t. If $d_t < d_s$, i.e., all AS paths over the satellite network are longer, then the local preference values of all AS paths in $\mathbf{P}^{Sat}_{AS_i}(NA_k)$ are set to $LocalPref(\mathbf{P}^{t}_{AS_i}(NA_k)) - 1$, i.e., $LocalPref(\mathbf{P}) = LocalPref(\mathbf{P}^{t}_{AS_i}(NA_k)) - 1$, $\forall \mathbf{P} \in \mathbf{P}^{Sat}_{AS_i}(NA_k)$.
7. If $d_t > d_s$, one of the AS paths over the satellite network is shorter, then the local preference values of all AS paths in $\mathbf{P}^{Sat}_{AS_i}(NA_k)$ except for $\mathbf{P}^{s}_{AS_i}(NA_k)$ are set to $LocalPref(\mathbf{P}^{t}_{AS_i}(NA_k)) - 1$, i.e., $LocalPref(\mathbf{P}) = LocalPref(\mathbf{P}^{t}_{AS_i}(NA_k)) - 1$, $\forall \mathbf{P} \in \mathbf{P}^{Sat}_{AS_i}(NA_k)$ and $\mathbf{P} \neq \mathbf{P}^{s}_{AS_i}(NA_k)$. The local preference value of $\mathbf{P}^{s}_{AS_i}(NA_k)$ is set to $LocalPref(\mathbf{P}^{t}_{AS_i}(NA_k)) + 1$.
8. The updated local preference values are advertised in the autonomous system AS_i.

4.3. BGP-S connection termination

Assume that a BGP-S connection between two peer gateways PGW_i and PGW_j is terminated because PGW_j does not receive any message from PGW_i within a time-out period. If the connection terminates due to time-out, PGW_j notifies the *APR* about the termination. *APR* checks if PGW_i is alive. If PGW_i is alive, no action is taken. If PGW_i does not respond, then *APR* records the failure of PGW_i in its database and informs all active peer gateways about this. Any existing connections to PGW_i is terminated and all RIB entries that use AS_i are withdrawn by active peer gateways within their ASs.

On the other hand, if a peer gateway PGW_i will be turned off or if AS_i does not want to receive any traffic from the satellite network, then PGW_i terminates all active connections with *NOTIFICATION* messages. The peer gateways that receive *NOTIFICATION* messages do not contact *APR*. Then PGW_i sends a message to *APR* indicating that it is no longer active. *APR* records this in its database and forwards this message to all active peer gateways. All RIB entries that use AS_i are withdrawn by active peer gateways.

5. Performance evaluation

We evaluated the performance of BGP-S on an integrated terrestrial/satellite IP networks model. This integrated network model consists of terrestrial ASs and a satellite network. The new network generation tool we use to create the integrated network topology is called the Integrated Terrestrial/Satellite Topology Generator (ITSTG). The performance of BGP-S is evaluated with simulations run on the network topologies created by ITSTG.

5.1. Topology generation process

The structure of ITSTG is shown in figure 4. It is extended from the Boston University Representative Internet Topology gEnerator (BRITE) [10]. The parameters for terrestrial and satellite networks are specified in a configuration file. Using these parameters, the topology generation engine generates terrestrial AS-level topology and the satellite network topology separately. Finally these two topologies are used as the inputs to the topology integrator to create the integrated terrestrial/satellite network topology. The terrestrial, satellite, and integrated terrestrial/satellite topologies are generated as described in the following sections.

5.1.1. Terrestrial topology

The specific details regarding how a terrestrial topology is generated depend on the specific generation model being used. In general, the generation process is divided into three steps:

1. *Place the nodes.* The nodes are placed on the terrestrial sphere with *heavy tailed* distribution, which describes the topological properties of the Internet [11]. The sphere is divided into squares, and each square is assigned a number of nodes drawn from a heavy tailed distribution. Then these nodes are placed randomly in the square. The positions of nodes have longitude in $[0°, 360°)$ and latitude in $[-90°, 90°]$.
2. *Interconnect the nodes.* The methods of interconnecting nodes are different for two different models: Waxman and Barabasi. The placing procedures are taken from BRITE with minor modifications.
   - *Waxman model.* In Waxman model [12], a new node tends to be connected to existing node(s) that are closer in distance. The nodes are added into the topology in an incremental way. A node is selected randomly to join the network and interconnected to other existing nodes. The incremental growth is a possible cause for power law of "outdegree exponent" [11] in any network topology.
   - *Barabasi model.* Barabasi model is proposed by Barabási and Albert [13]. This model suggests two possible causes for power law of "outdegree exponent" in network topologies: "incremental growth" and "preferential connectivity". Incremental growth refers to growing networks that are formed by the continuous addition of new nodes, which simulates the gradual increase in the size of the network. Preferential connectivity refers to the tendency of a new node to connect to existing nodes that already have high connectivity.
3. *Assign attributes.* The bandwidth of a terrestrial link is assigned a value randomly drawn between *BWmin* and *BWmax*. The delay within an AS is a random variable uniformly distributed within *ASmin* and *ASmax*, which are specified in the configuration file.

5.1.2. Satellite topology

We consider a single-layer Walker Star [14] type LEO satellite network. Satellites are placed on the sphere of radius $R_E + h$, where h is the altitude of LEO satellites. We utilize the "logical location" concept in [3]. The *logical locations* are equally spaced points in the grid of the LEO satellite constellation. They do not move with respect to the Earth and are embodied by the nearest LEO satellites. The communication between the satellites occurs through inter-satellite links (ISLs). The neighboring satellites within the same plane are

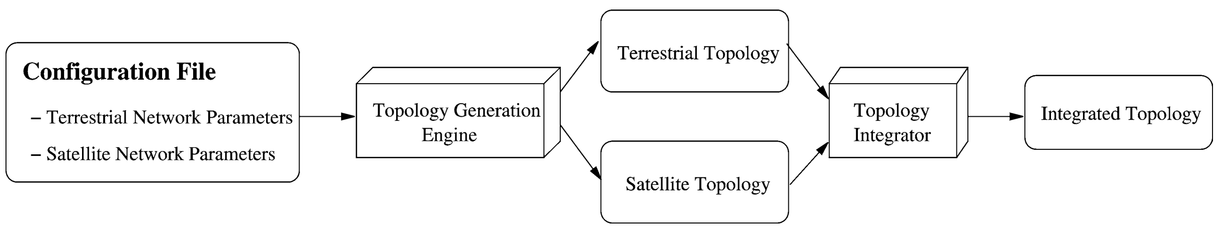

Figure 4. The schematic structure of ITSTG.

connected via intra-orbit ISLs. Similarly, the neighbors in adjacent planes are connected through inter-orbit ISLs. One satellite can have at most four adjacent links. Inter-orbital links only exist between neighboring satellites outside polar areas. The bandwidth of links in the satellite network is fixed.

5.1.3. Integrated terrestrial/satellite topology

In BGP-S, the detailed topology of the satellite network is hidden from the terrestrial network. The communication between satellite and terrestrial parts of this integrated topology is accomplished through the terrestrial gateways. We assume that one gateway belongs to one terrestrial AS and has only one user data link (UDL) to one satellite, which is represented by the nearest logical location.

The generation process of integrated network topology has the following three steps:

1. *Interconnect gateways and satellites.* First, we select the value of p, the percentage of ASs having connections with the satellite network. For every terrestrial node, a value is randomly generated between 0 and 1. If the value is smaller than p, then a link is added between that node and its nearest satellite logical location.
2. *Condense the satellite topology.* The condensed satellite topology only includes the satellites that have UDLs. These satellites are referred to as representative nodes (RNs). Virtual links are built between every pair of representative nodes. The cost of the virtual link between two representative nodes RN1 and RN2 is the accumulated cost of the nodes and links along the path from RN1 to RN2 within the satellite network. If we define the cost as the delay of the link, then the cost of a virtual link is the sum of delays of all links along the minimum delay path. A virtual link is counted as one hop.
3. *Create the integrated topology.* The integrated topology is the combination of terrestrial topology and the condensed satellite topology. It includes all terrestrial nodes and links, the satellite representative nodes and links between them, and the links between terrestrial gateways and satellite representative nodes.

Figure 5(a) gives an example of interconnected terrestrial/satellite network, in which the terrestrial network has five ASs, three of which contain gateways. After going through the above three steps, the integrated terrestrial/satellite topology is generated as shown in figure 5(b). In the satellite part, only the three satellites that have UDL connections to terrestrial gateways are kept in the condensed topology. The dashed lines are virtual links that connect these three satellites.

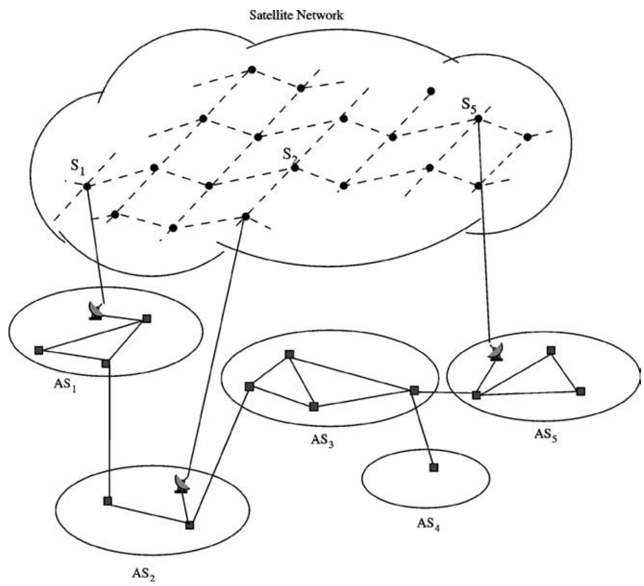

(a) Interconnection of networks.

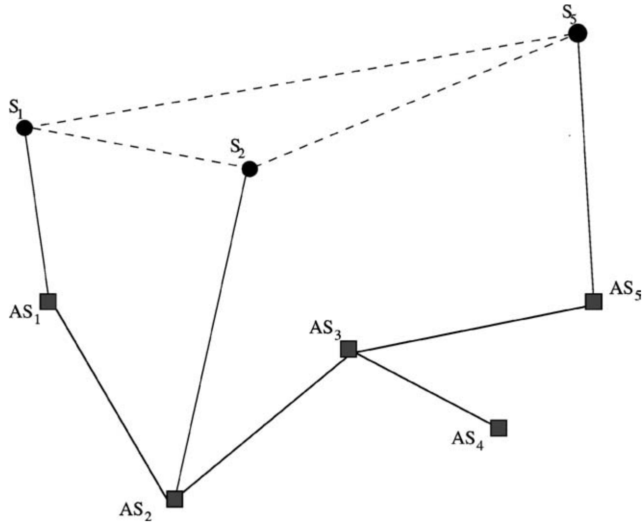

(b) Condensed topology.

Figure 5. Integrated terrestrial/satellite topology generation.

5.2. Simulation results

Based on the topology created, we have simulated the routing between any two terrestrial AS nodes. In the simulations, we do not consider the source or destination located in the satellite network, as the users and service providers reside on Earth. In BGP-4 protocol, the configured policies override the efficiency considerations in the path selection process [8]. We cannot simulate BGP-4 as it works in the real Internet because it is not possible to make realistic assumptions about the administrators' preferences. Hence, we implement minimum hop routing to reflect the characteristics when AS-path hop length is the decision criterion of choosing the path in BGP-4.

In our integrated terrestrial/satellite network, every link is associated with an instantaneous delay. This link can be an intra-AS link on earth, a UDL between a gateway and a satellite, or an ISL. Each link is modeled as an infinite capacity queue. Given link load and link capacity, with the assumption of Poisson arrival rate and exponentially distributed service time, the queuing delay of each link can be deduced by the $M/M/1$ queuing model. As the terrestrial part of the ITSTG is built on AS-level, a packet also experiences delay within an AS, which is represented by an intra-AS delay value of the AS node.

Table 1
Simulation parameters.

Terrestrial	Satellite
$LS = 10$	Number of plane = 12
$m = 2$	Number of satellites per plane = 24
Link bandwidth	ISL bandwidth = 160 Mbps
(BWmin = 10 Mbps, BWmax = 1 Gbps)	
Intra-AS delay	Altitude = 1400 km
(ASmin = 5 ms, ASmax = 50 ms)	

In all simulations, the number of nodes in the terrestrial network is chosen as 3000. The simulation parameters are listed in table 1. In this table, $LS = 10$ stands for a side length of 10° for the square used in the *heavy tailed* node-placement method, and m stands for the number of links per new node. The bandwidths of inter-AS links and the intra-AS delays are uniformly distributed between selected minimum and maximum values. The link loads are uniformly distributed between 0% and 100% of the respective link bandwidths. The bandwidth of UDLs is set as 1.6 Mbps.

We conducted simulations on the integrated topology and compared the delay performance of different routing policies. The terrestrial part is generated from Waxman and Barabasi models. For different values of percentage p, which is the ratio of AS nodes having satellite connections, we generated 100 different integrated terrestrial/satellite network topologies. Taking a topology generated independently each time, we chose 100 different source–destination pairs. For each source–destination pair, BGP-S and BGP-4 are run separately, and the delay results for different routing method are recorded. This procedure is repeated for every topology. The delay comparisons are made by averaging all $100 \cdot 100 = 10{,}000$ results.

Performance comparison of BGP-4 in integrated and terrestrial networks. The first set of the simulations compares the delay metric when BGP-4 is implemented both in the terrestrial network and the integrated terrestrial/satellite network. Figure 6 gives the ratio of path delay by implementing BGP-4 globally with and without satellite network versus p. If the ratio equals 1, it means that the delay of the path selected by BGP-4 does not change after the satellite network is included. If the ratio is larger than 1, it means that including satellite network in route selection of BGP-4 introduces longer delays. If the ratio is less than 1, the path delay will be reduced if satellite AS is included in BGP-4. Figure 6 shows that for Waxman model and Barabasi model (when p is larger than 15%), if we apply BGP-4 in the integrated terrestrial/satellite network, the delay is smaller than that in terrestrial network alone. This shows that when satellite links are included in routing selection, the performance improves in terms of delay metric. As p increases, it is easier for BGP-4 to choose the path through satellite network, the performance of BGP-4 in the terrestrial/satellite network gets better.

Performance comparison between BGP-S and BGP-4. For the following set of experiments, BGP-S and BGP-4 are simulated on the integrated terrestrial/satellite network. Their results are compared according to the path delay characteristics. The ratio of the delays using BGP-S and BGP-4 versus p is depicted in figure 7. If the ratio is larger than 1, it means that implementing BGP-S increases the delay of path. Otherwise, the path delay will be reduced if BGP-S is used. In this figure, the ratio is always less than 1 for both Waxman and Barabasi models, which means that BGP-S always produces lower delays than BGP-4 in the integrated network. This set of experiments show that the satellite network can be utilized with BGP-S in a much better way. However, the decrease/increase of delay ratio as p increases depends on the specific model (Waxman or Barabasi) used for the terrestrial AS-level network. Later simulations will show that the change also varies with different selections of satellite constellation.

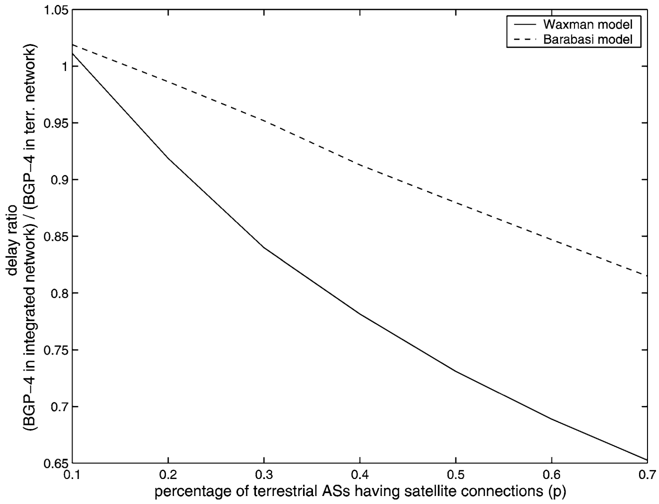

Figure 6. Performance comparison of BGP-4 in integrated and terrestrial networks.

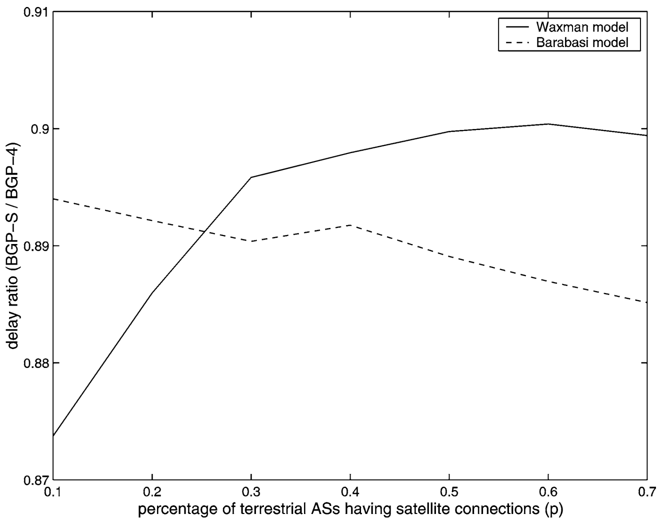

Figure 7. Performance comparison between BGP-S and BGP-4.

Effect of satellite parameters on BGP-S performance. The performance of BGP-S is affected by the architecture of satellite network, such as the number of nodes in satellite network

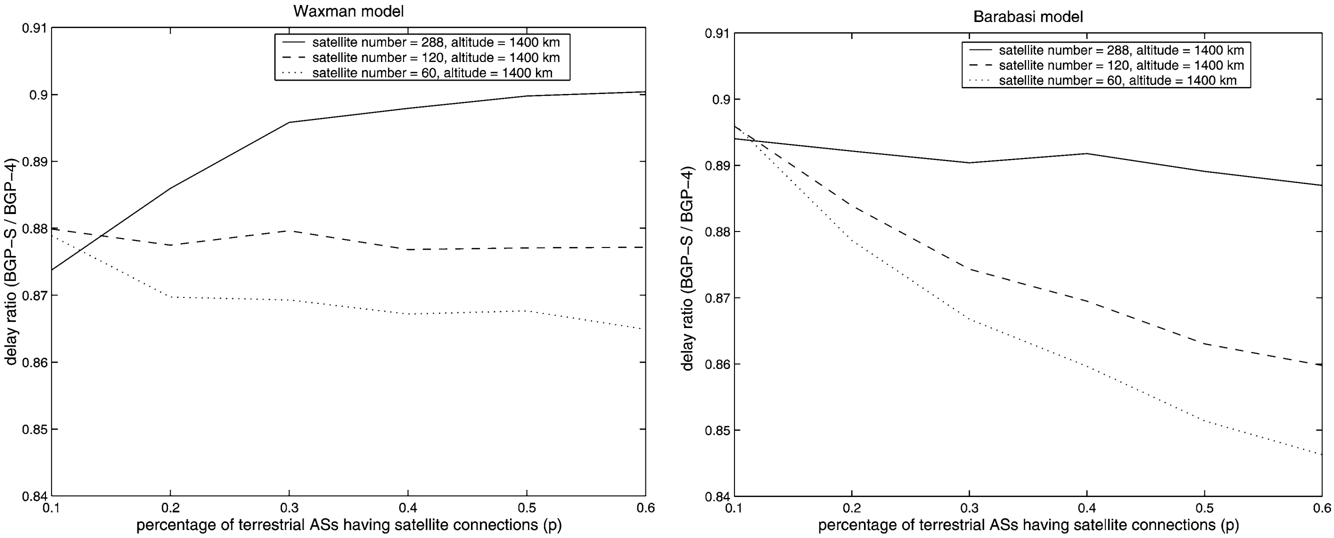

(a) Different satellite numbers.

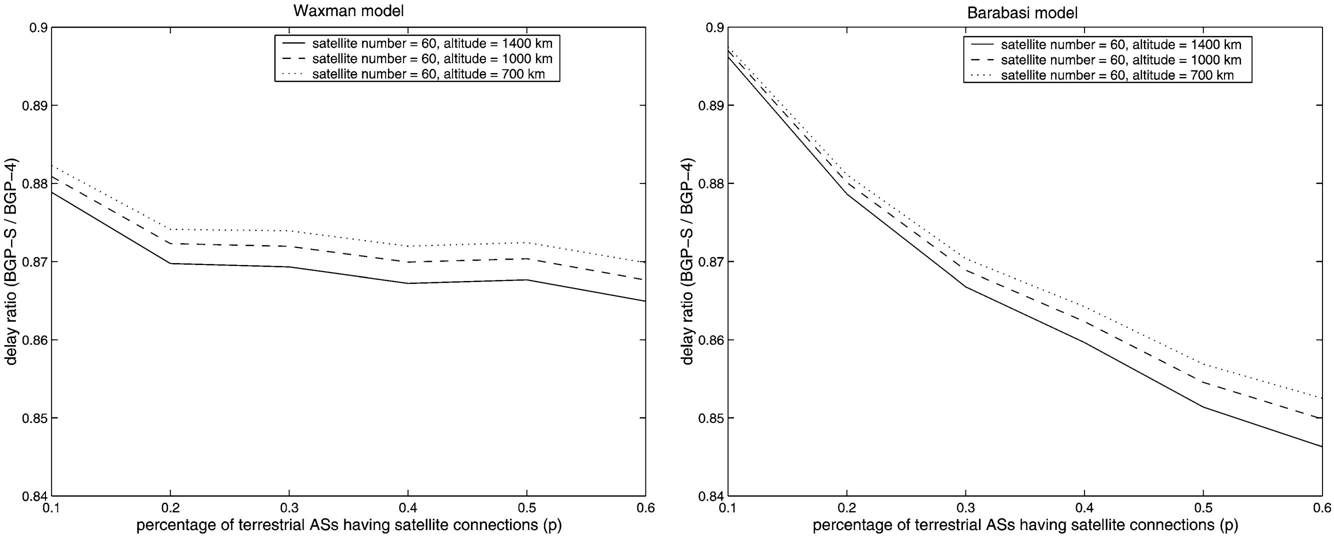

(b) Different satellite altitudes.

Figure 8. Effect of satellite parameters on BGP-S performance.

and the altitude of the satellite layer. In this set of experiments, we show the effect of satellite network architecture on the BGP-S performance. In figure 8(a), the delay ratio of BGP-S and BGP-4 versus p is depicted for satellite architectures with different satellite numbers. The altitudes of all three architectures are fixed as 1400 km, whereas the satellite numbers are chosen as 60 (with 6 planes), 120 (with 10 planes) and 288 (with 12 planes), respectively. It shows that when the satellite number decreases, the delay ratio is smaller. In fact, BGP-S produces similar results in all three different architectures. However, as the satellite number decreases, the paths selected by BGP-4 give longer delay as the satellite nodes become sparse.

Next, we fixed the satellite number as 60, and changed the altitude of satellite layer as 700 km, 1000 km, and 1400 km. The routing procedure is repeated for all three architecture independently. The delay ratio of BGP-S and BGP-4 in the integrated satellite/terrestrial network versus p is plotted in figure 8(b). It can be seen that as the altitude of satellite layer increases, the performance of BGP-S gets better. This is because when the altitude of satellites is higher, the hops represented by UDLs to/from satellites are longer. If BGP-4 chooses such links, the selected path introduces longer delay. However, BGP-S also gives longer delay as the satellite altitude grows, as the result, the delay ratio varies only slightly (within 1%) under the three different architectures. Hence, we conclude that the satellite altitude does not affect much on the performance gain of BGP-S over BGP-4.

Effect of gateway selection methods on BGP-S performance. In previous experiments, the peer gateways are positioned randomly according to the explanation in section 5.1.3. However, we expect that some AS nodes are more likely to have connections to the satellite network. These nodes may include the backbone nodes (e.g., Tier-1 ISPs) and remote nodes (e.g., stub ASs which are several hops away from the Tier-1 ISPs). Hence, in this section, another method called "*filtered gateway selection*" is used to place the gateways. We set $m = 1$ in

Table 2
Values of p under filtered gateway selection method.

	n_r							
	1	2	3	4	5	6	7	8
Waxman model ($n_b = 9$)	0.506	0.497	0.470	0.416	0.340	0.254	0.173	0.109
Barabasi model ($n_b = 25$)	0.670	0.591	0.432	0.260	0.132	0.059	0.024	0.010

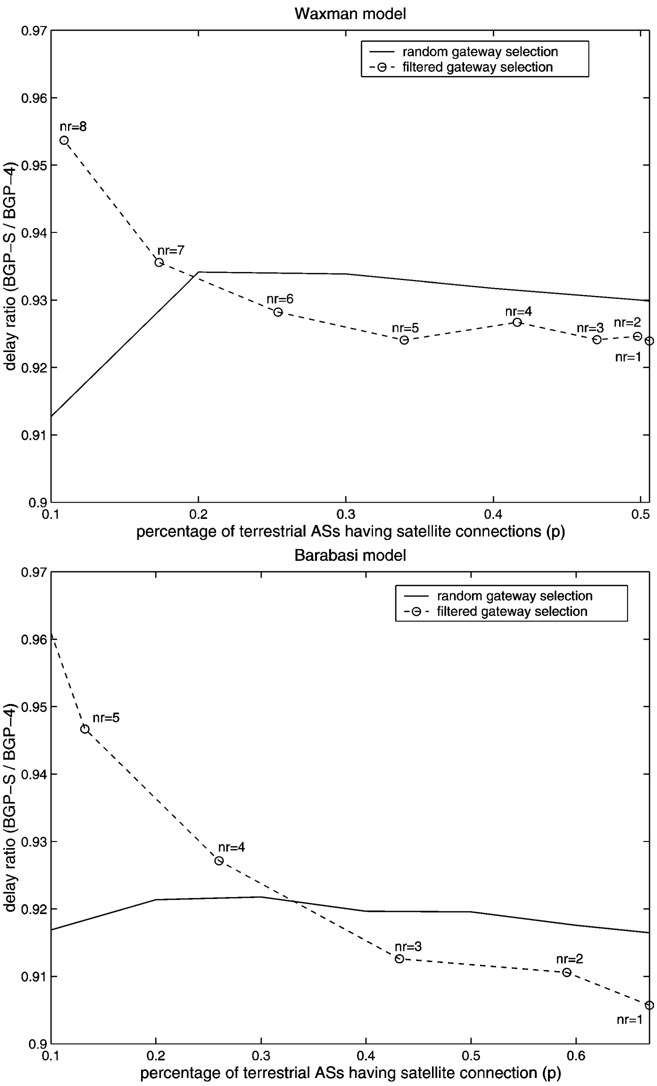

Figure 9. Effect of gateway selection methods on BGP-S performance.

this set of experiment, thus, the stub ASs are those with node degree equal to 1.

First, we search for the backbone nodes and remote nodes, where backbone nodes are those with outdegree larger or equal to n_b, remote nodes are the terrestrial nodes with node degree equal to 1 and are n_r hops away from all the backbone nodes. If a node is either a backbone node or a remote node, a peer gateway is equipped and a link is added between this node and its nearest satellite logical location. After placing all gateways, the satellite topology is condensed and the integrated topology is created. We chose several different values of n_b and n_r for Waxman model and Barabasi model, respectively. This gave different percentage p of AS nodes having satellite connections. Table 2 lists the values of p corresponding to different n_b and n_r. Because the topologies built by the Waxman and Barabasi models have different node degree distributions, their n_b values are different, the change of n_r also maps to different values of p for the two models.

Figure 9 shows the performance comparison of BGP-S under *random gateway selection* method and *filtered gateway selection* method. The y-axis represents the ratio of BGP-S and BGP-4 in the integrated terrestrial/satellite network. Note that the results for random gateway selection method are different from those in figure 7 due to different m values. The results show that as p grows, BGP-S performs better under filtered gateway selection method. Moreover, the growth of p has greater effect on the delay ratio of BGP-S and BGP-4 in the integrated terrestrial/satellite network, if the gateways are installed selectively in backbone ASs and remote ASs.

6. Conclusion

In this work, we introduced the BGP-S protocol for integration of the IP-based satellite networks with the Internet. The BGP-S protocol does not require a special satellite network architecture and works independent of the internal routing of the satellite network. BGP-S is fully compatible with the BGP-4 protocol. It uses BGP-4's *LocalPref* value to propagate the paths it learned through the satellite network. To accept or reject the paths learned through the satellite network, the policies manually configured in the routers are given priority. If alternative paths are available, the choice is based on measuring the delays on the existing paths. BGP-4 policies are directly adopted for full compatibility with existing systems. Moreover, BGP-S is implemented only in one terrestrial gateway in every terrestrial AS to reduce the complexity. The functionalities of the BGP speakers in the terrestrial ASs remain the same. The performance of BGP-S has been assessed with simulations. The results show that BGP-S always produces lower delays than BGP-4 in the integrated terrestrial/satellite network. When satellite altitude increases or satellite number decreases, the performance gain of BGP-S over BGP-4 grows. The effect of terrestrial gateway selection method on BGP-S performance is also evaluated.

Acknowledgements

This work is supported by the National Science Foundation under Grant ANI-0087762 and NASA Glenn Research Center under Grant NAG3-2580.

References

[1] Y. Hu and V.O.K. Li, Satellite-based Internet: A tutorial, IEEE Communications Magazine 39 (March 2001) 154–162.

[2] Y. Hashimoto and B. Sarikaya, Design of IP-based routing in a LEO satellite network, in: *Proceedings of 3rd ACM/IEEE International Workshop on Satellite-based Information Services (WOSBIS'98)* (October 1998) pp. 81–88.

[3] E. Ekici, I.F. Akyildiz and M.D. Bender, A distributed routing algorithm for datagram traffic in LEO satellite networks, IEEE/ACM Transactions on Networking 9(2) (2001) 137–147.
[4] T.R. Henderson and R.H. Katz, On distributed, geographic-based packet routing for LEO satellite networks, in: *Proceedings of GLOBECOM 2000*, Vol. 2 (December 2000) pp. 1119–1123.
[5] C. Chen, E. Ekici and I.F. Akyildiz, Satellite grouping and routing protocol for LEO/MEO satellite IP networks, in: *Proceedings of the 5th ACM International Workshop on Wireless Mobile Multimedia (WoWMoM'02)* (September 2002) pp. 109–116.
[6] P. Narvaez, A. Clerget and W. Dabbous, Internet routing over LEO satellite constellations, in: *Proceedings of 3rd ACM/IEEE International Workshop on Satellite-Based Information Services (WOSBIS'98)* (October 1998).
[7] L. Wood, A. Clerget, I. Andrikopoulos, G. Pavlou and W. Dabbous, IP routing issues in satellite constellation networks, International Journal of Satellite Communications 19 (January/February 2001) 69–92.
[8] Y. Rekhter and T. Li, A Border Gateway Protocol (BGP-4), RFC 1771 (March 1995).
[9] J.W. Stewart III, *BGP-4: Inter-Domain Routing in the Internet* (Addison-Wesley, 1999).
[10] BRITE, BRITE: Boston university representative internet topology generator, `http://www.cs.bu.edu/brite/index.html`
[11] M. Faloutsos, P. Faloutsos and C. Faloutsos, On power-law relationships of the Internet topology, ACM Computer Communication Review, Cambridge, MA (September 1999).
[12] B.M. Waxman, Routing of multipoint connections, IEEE Journal of Selected Areas in Communications 6(9) (December 1998) 1617–1622.
[13] A.L. Barabási and R. Albert, Emergence of scaling in random networks, Science (October 1999) 509–512.
[14] C.J. Wang, Structural properties of a low earth orbit satellite constellation – the Walker Delta Network, in: *Proceedings of MILCOM'93*, Vol. 3 (1993) pp. 968–972.

Eylem Ekici has received his B.S. and M.S. degrees in computer engineering from Bogazici University, Istanbul, Turkey, in 1997 and 1998, respectively. He received his Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, in 2002. In October 2002, he joined the Department of Electrical Engineering of the Ohio State University, Columbus, OH.

Dr. Ekici's research interests include space-based networks, wireless sensor networks, and next generation wireless networks, with a focus on multi-access control, routing and multicasting protocols, and resource management. He also conducts research on interfacing dissimilar networks.

E-mail: ekici@ece.osu.edu

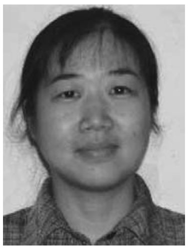

Chao Chen received the B.S. and M.S. degrees from Deparment of Electronic Engineering, Shanghai Jiao Tong University, Shanghai, China in 1998 and 2001, respectively. She is currently working toward the Ph.D. degree in the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta. She is a graduate research assistant in the Broadband and Wireless Networking Laboratory at the Georgia Institute of Technology. Her research interests include routing in satellite networks, wireless networks, and interplanetary networks.

E-mail: cchen@ece.gatech.edu