BORDER GATEWAY PROTOCOL 4 (BGP4) PERFORMANCE OVER INTERMITTENT SATELLITE LINKS**

R.G. Cole , L. Benmohamed, A. DeSimone and B. Doshi*

Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Road, Laurel, MD, 20723

ABSTRACT

We performed an investigation of the Border Gateway Protocol 4 (BGP-4) [9] behavior in the presence of intermittent satellite channels. Specifically, we developed an analytic model of BGP-4 state transitions and use this model to estimate bandwidth consumption of BGP-4 over the intermittent channels. We compare our model's estimates with fidelity simulation results of bandwidth consumption of BGP-4. We briefly discuss this work in the context of the Global Information Grid (GIG) and our ongoing efforts to develop a system wide simulation model of the GIG.

1. INTRODUCTION

The U.S. Government is building the Global Information Grid (GIG), which will support convergence of its current disparate communication networks onto a common IP network. A critical aspect of its design is the inter-domain BGP-4 routing architecture required to meet the diverse requirements imposed by its various user networks. BGP-4 will be required to handle routing through a vast array of networking technologies, including intermittent satellite links, high speed fiber backbones, highly mobile networks and numerous types of wireless links.

In this paper, we investigate the performance and behavior of BGP-4 in the presence of intermittent links. We develop analytic models for engineering studies and compare with simulation models of BGP-4. There exist few, high fidelity simulation studies which model BGP-4 performance [8] [5]. To our knowledge, no studies examine the impact of running BGP-4 over intermittent satellite links.

2. BGP-4 ANALYSIS

Various types of satellite and terminal configurations are deployed or planned for the GIG. BGP-4 must support routing over these intermittent links. These satellite connections suffer channel fading due to obstructions, rain, antenna placement on ships and other mobile platforms. BGP-4 must support routing over intermittent links, requiring the proper protocol tuning and local decision policies. Here, we develop an analytic model of the state of the BGP-4 link running over an intermittent channel. We use this BGP-4 state model to derive the bandwidth required to support BGP-4 status, policies and routing updates.

When two peer BGP-4 routers first come up, they enter an ACTIVE state [9]. In the ACTIVE state, they frequently attempt to establish a TCP connection between themselves. Once a TCP connection is established, the BGP-4 ESTABLISHED state is entered. The peers then negotiate protocol parameters, and transmit their entire network layer reachability information, or prefixes. In the ESTABLISHED state, the BGP-4 routers sense the state of the underlying communications channel with the BGP-4 KEEP ALIVE timer mechanism. Upon the expiration of the KEEP ALIVE timer, the BGP-4 router transmits a keep alive message to its peer if no other BGP-4 traffic had been sent over the previous keep alive time interval. If the peer router fails to receive any BGP-4 traffic over the channel within the HOLD DOWN time interval, it terminates the TCP connection and transitions back to the ACTIVE state.

To derive our model, we first consider a relatively simple Gilbert model of the satellite channel [3]. The satellite channel is characterized by two states, i.e., a good state with zero packet loss and a bad state with a probability of packet loss of unity. BGP-4 periodically samples the existence of the links with its peers through its KEEP ALIVE mechanism. Hence, we convert the channel model to a BGP-4-sensed channel model, where the sensing is performed by BGP-4's KEEP ALIVE mechanism. We then derive expressions for the frequency of BGP-4 routing updates over the channel by treating the BGP-4 peering state as the result of applying a low band-pass filter to the BGP-4-sensed channel state. The characteristic of the filter is directly related to the BGP-4 HOLD DOWN timer. Our process for modeling this behavior is illustrated in Figure 1. The result is a model of the frequency of BGP-4 state thrashing over an intermittent communications channel as a function of channel characteristics and protocol parame-

---

* R.G. Cole is the corresponding author, email: robert.cole@jhuapl.edu.

** This material is based upon work supported by the Defense Advanced Research Projects Agency Information Processing Technology Office Network Modeling and Simulation Program. Issued by DARPA/CMO under Contract No. MDA 972-01-D-0005, Task Order Number 0037.

ter values. This thrashing frequency is related to the bandwidth requirements of BGP-4 necessary to communicate reachability information.

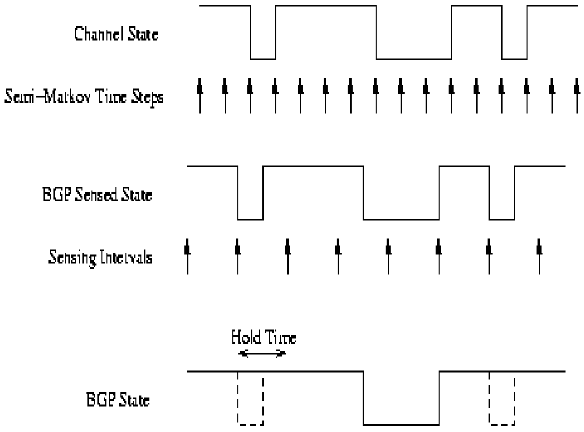

Fig. 1. The model development process.

Our Gilbert model is characterized by δ, the number of seconds for which the system stays in a given state prior to making a transition decision, and the transition probabilities between the good-to-good and the bad-to-bad states, p_{gg} and p_{bb} respectively. We assume that the underlying Gilbert model describes the actual state of the channel. The "Sensed-State" of the link is also a Two-State, Gilbert Model where we assume that the BGP-4 peer senses the underlying channel each M time increments, δ, i.e., $M\delta$ equals the BGP-4 KEEP ALIVE timer. Hence, we define the primed transition probabilities as the probabilities that, in M transitions, the sensed-state transitions from good-to-good or bad-to-bad based upon the following iterative equations (assuming $M > 1$)

$$p'_{gg}(M) = p_{gg}p'_{gg}(M-1) + p_{gb}p'_{bg}(M-1) \tag{1}$$

and

$$p'_{bb}(M) = p_{bb}p'_{bb}(M-1) + p_{bg}p'_{gb}(M-1) \tag{2}$$

We transform these iterative equations to

$$\begin{aligned} p'_{gg}(M) &= p_{gg}p'_{gg}(M-1) \\ &\quad +(1-p_{gg})(1-p'_{bb}(M-1)) \end{aligned} \tag{3}$$

$$\begin{aligned} p'_{bb}(M) &= p_{bb}p'_{bb}(M-1) \\ &\quad +(1-p_{bb})(1-p'_{gg}(M-1)) \end{aligned} \tag{4}$$

where $p'_{gg}(1) = p_{gg}$ and $p'_{bb}(1) = p_{bb}$. To complete our Sensed-State Model we replace the time increment, δ, with $M\delta$.

To estimate the BGP-4 traffic over the satellite channel, we need to compute the mean time in the BGP-4 ESTABLISHED state given a HOLD DOWN timer of $NM\delta$ seconds. We denote this mean time in the good and bad states as $t_g^{(N)}$ and $t_b^{(N)}$ respectively. We compute $t_g^{(N)}$ using the following recursive expression

$$\begin{aligned} t_g^{(N)} &= 1 + p'_{gg}t_g^{(N)} \\ &\quad +p'_{gb}\sum_{i=1}^{N}(i+t_g^{(N)})Pr\{t_B = i\} \\ &\quad +p'_{gb}NPr\{t_B > N\} \end{aligned} \tag{5}$$

which we solve to yield

$$t_g^{(N)} = t_g^{(0)}\left[\frac{1+p'_{gb}f^{(N)}(p'_{bb})}{1-h^{(N)}(p'_{bb})}\right] \tag{6}$$

where

$$\begin{aligned} f^{(N)}(x) &= N(1-x)x^{(N+1)} \\ &\quad +\frac{1-(N-1)x^N + Nx^{(N+1)}}{(1-x)} \end{aligned} \tag{7}$$

$$h^{(N)}(x) = 1 - x^N \tag{8}$$

and $t_g^{(0)} = (1-p'_{gg})^{-1}$. It turns out that $t_b^{(N)} = t_b^{(0)}$, due to the memoryless nature of the Markov model and the mechanisms in BGP-4 for transitioning from the ACTIVE and ESTABLISHED state. Thus, our expression for the BGP-4 overhead on an intermittent satellite link is

$$T = \frac{\alpha}{t_{KEEPALIVE}} + \frac{\beta}{t_{cycle}} \tag{9}$$

where α and β are fitting parameters, $t_{KEEPALIVE}$ is equal to the KEEP ALIVE timer value, and

$$t_{cycle}^{(N)} = t_g^{(N)} + t_b^{(0)} \tag{10}$$

The first term in Eq.(9) reflects the chatter overhead required to maintain the peering session and the second term is the overhead of advertising prefixes over the intermittent link. The overhead includes all BGP-4 messages and communications related to the underlying TCP connection. The term, β, includes control traffic related to prefix advertisements and initial protocol negotiations after transitioning to the ESTABLISHED state. Ignoring the traffic due to negotiation and drawing upon [10], we can interpret β as

$$\beta = (N + MA)O_L \tag{11}$$

where N is the number of prefixes advertised within the entire network, A is the number of ASes in the network, M is the average number of AS hops in a prefix AS-PATH attribute, and O_L is the protocol overhead in advertising the prefixes and associated attributes. Then, Eqs. (10) and (11) give the expected BGP-4 traffic over an intermittent link in terms of channel and protocol parameters. These estimates represent a conservative prediction as techniques exist, e.g., route aggregation, which can reduce the size of router advertisements under reasonable topology assumptions, as discussed in [9].

3. BGP-4 SIMULATIONS

We have developed a scalable and easily parameterized model of the GIG topologies for numerous performance studies, see [1]. We use the Network Simulator 2 (NS2) [6] and BGP++ [2] simulation tools, which require configuration in TCL and GNU-Zebra style BGP-4 router configurations [4]. Using BGP++ offers us a high fidelity representation of the BGP-4 protocol mechanisms and protocol exchanges over the simulated network.

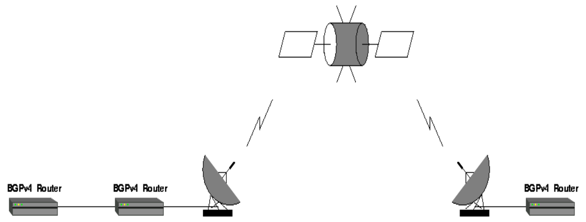

Fig. 2. The BGP-4 reference satellite configuration.

To test our model, we made a series of simulation runs of a simple three router network. Figure 2 shows the simulated configuration. It consists of three BGP-4 routers sitting in three distinct Autonomous System (AS) domains. The two routers on the Left Hand Side (LHS) of the figure are peering over a T1 wired facility, while the two routers on the Right Hand Side (RHS) of the figure are peering over a 1.5 Mbps satellite link. Each router represents a single AS advertising 1000 prefixes. We assume that all prefixes are advertised over the link, i.e., that no BGP-4 policy is applied. We set the HOLD DOWN timer to equal three times the KEEP ALIVE timer, as recommended in [9]. We set the transition parameters for the good and bad states of the satellite channel to be identical. The mean lifetime in the good state was set to 250, 500 and 1000 seconds, respectively. We ran each combination of parameters 200 times for a simulated duration of 20,000 seconds. The traffic metric is the mean number of bytes per second related to BGP-4 exchanges transmitted over the satellite link averaged over the length of each simulation run.

Figure 3 shows the results of the simulation runs. We see, as expected, that the mean traffic overhead decreases as the HOLD DOWN timer increases due to both the lessening of the effects of the KEEP ALIVE chatter and the prefix advertisements over the intermittent links. Also, as the mean lifetime of the good state increases, the amount of traffic overhead decreases.

We performed a relatively simple fitting procedure to test our expression, Eq.(9). We first set $\beta = 0$ and varied α until we fit reasonably well the maximum traffic point in Figure 3. This yielded $\alpha = 23,000$ and is shown in Figure 5. As the α parameter is related only to the HOLD DOWN timer, we see that all curves for different satellite lifetimes overlap. We then set $\alpha = 0$ and varied β to fit the maxi-

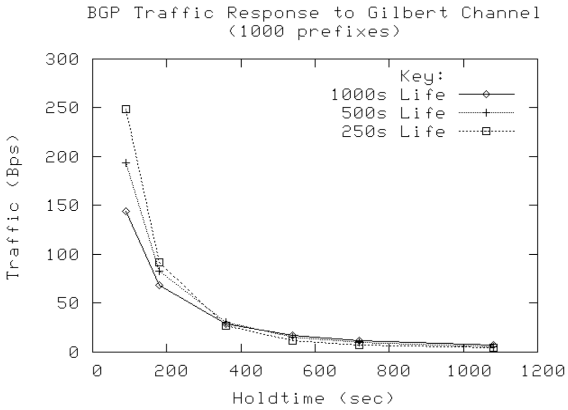

Fig. 3. The simulation data of BGP-4 traffic.

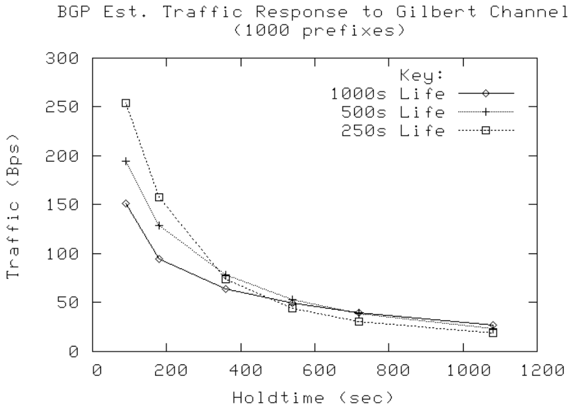

Fig. 4. The fit to the simulation data.

mum traffic point in Figure 3. This yielded $\beta = 180,000$ as shown in Figure 6. We then took a weighted average of these two cases where the weightings are described by $\gamma\alpha$ and $(1-\gamma)\beta$ and found that $\gamma = 3/8$ gave a reasonable fit. This final case is shown in Figure 4 and is to be compared with Figure 3. We see that this expression (with only two free parameters) does a reasonable job of characterizing the dependence of the BGP-4 traffic overhead on the HOLD DOWN and KEEP ALIVE timers and the satellite channel model. This expression also shows a cross-over behavior between the different satellite models as the HOLD DOWN timer is increased from 90 seconds up to 1080 seconds as found in the simulation runs.

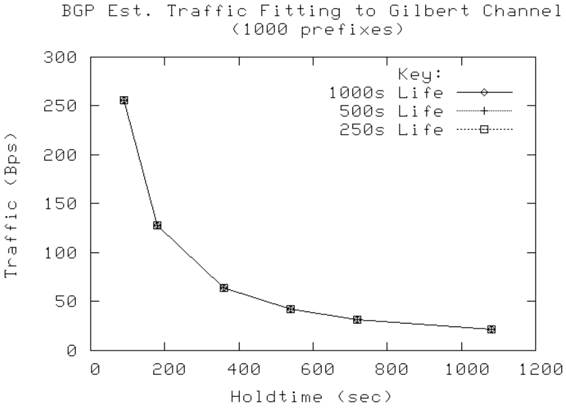

Fig. 5. The fitting to the simulation data setting $\beta = 0$.

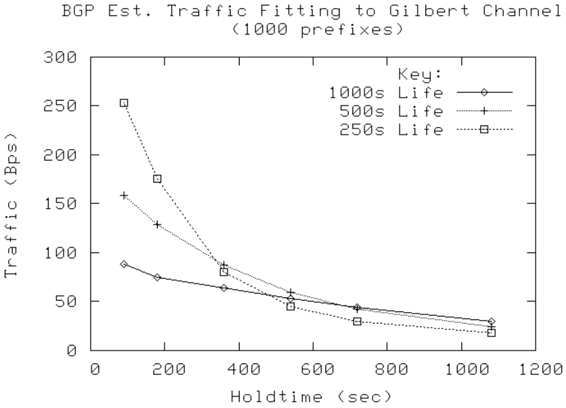

Fig. 6. The fitting to the simulation data setting $\alpha = 0$.

4. CONCLUSIONS

The GIG is planning on using BGP-4 as its inter-domain routing protocol. As such, we have embarked upon a series of studies investigating performance, scale and security issues relating to BGP-4 deployment in the GIG. In this initial study, we focused on the impact of running BGP-4 over an intermittent satellite link. Satellites are relied upon extensively by the DoD for connectivity to remote and increasingly to mobile platforms. As an aid to the design, configuration and deployment of BGP-4 connectivity over satellites, we developed an analytic model to predict the traffic overhead of running BGP-4 over a satellite channel. We validated the model against high fidelity simulation studies. We are in the process of running further BGP-4 routing studies for the Global Information Grid [1].

5. ACKNOWLEDGMENTS

We wish to thank Phil Chimento of the JHU Applied Physics Laboratory for numerous and helpful discussions on this and related work. We also wish to gratefully acknowledge Xenofontas Dimitropoulos of the Georgia Tech University for his support in answering our questions regarding the use and configuration of the BGP++ simulation tool.

6. REFERENCES

[1] Cole, R.G., Benmohamed, L., Chimento, P., DeSimone, A. and B. Doshi, *Initial Investigations of BGP Performance on a Global Information Grid Simulation Platform*, submitted to MILCOM 2005, Atlantic City, NJ, 2005.

[2] Dimitropoulos, X.C., and G. Riley, *Simulation Tool for BGP Studies*, *http://www.ece.gatech.edu/research/labs/ MANIACS/BGP++/*, 2005.

[3] Gilbert, E.N., *Capacity of a burst-noise channel*, Bell Syst. Tech. J., vol.39, pp 1253-1265, September, 1960.

[4] GNU Zebra, *The GNU Zebra BGP Daemon*, http://www.gnuzebra.org/, 2005.

[5] Liu, X. and A.A. Chien, *Realistic Large-Scale Online Network Simulation*, preprint, 2004.

[6] Network Simulator 2, *http://www.isi.edu/ns2/nam/*, 2003.

[7] Nichols, D., et al., *SSFNet*, *http://www.ssfnet.org/*, 2005.

[8] Quoitin, B., et al., *A performance evaluation of BGP-based traffic engineering*, preprint, 2004.

[9] Rekhter, Y. and T. Li, *The Border Gateway Protocol - version 4 (BGPv4)*, IETF RFC 1771, 1995.

[10] Traina, P., *BGP-4 Protocol Analysis*, IETF RFC 1774, 1995.