Building a Secure and Reliable Network via Multi-homed VPN

Baijian Yang
Department of Technology
Ball State University
byang@bsu.edu

Tianguang Gao
Department of Educational Technology
Ball State University
tgao@bsu.edu

Abstract

When it comes to network design, people are constantly looking for solutions to build more reliable and more secure networks at affordable costs. In this paper, we present our solution that can achieve the above goals by supporting Virtual Private Networks (VPN) over multi-homed networks. While the conventional VPN technology offers a cost-effective way to securely communicate through unsecured public networks, it is not reliable in the sense that when a single node or a link on the network path fails, the VPN will fall short. Another interesting technology, termed as Multi-homed network, can enhance the network availability by having more than one external links to the internet. Multi-homed networks are more reliable especially if the external links are offered by different Internet Service Providers (ISPs). When one network path from an ISP becomes unavailable, it can still talk to the outside world from another ISP. This article discusses and analyzes various issues when incorporating VPN with multi-homed networks. Issues like IPSec packets that will not travel through NAT are addressed and a number of workarounds are hence proposed and studied. The preliminary simulation results demonstrate our approach can provide a more reliable private network at much reduced costs.

Motivation

As the world has entered the digital era, having secure and reliable computer networks is vital to many enterprises to offer business transactions online, to allow its employees connect to its main servers through public, unreliable networks. While Virtual Private Networks (VPNs) technology has widely adopted in practices to make data transmission secure, it still suffers from network failures and congestions. This paper aims to solve the problem by implementing VPNs over multi-homed networks, i.e. networks have multiple external links so that VPNs can transparently switch from a congested or a downed public network to other public networks to improve reliability and load balancing.

Several proposals have been presented to make current internet connections more reliable. One of the techniques is using overlay network, which tries to build multiple links from the source to the destination instead of a single end-to-end link [5]. Therefore, when some of the nodes or links on the routing path are either congested or downed, the overlay networks can simply deliver data using other paths. The advantage is obvious: the end users do not have to wait for the slow convergence of network and thus has the potential to benefit both consumers and the internet service providers [4]. A typical way to implement overlay network is to modify the current design of Border Gate Protocol (BGP) such that the edge routers of each domain can keep the status of multiple links. However, the routing entries of overlay networks are very hard to be integrated into the current existing classless inter-domain routing scheme and therefore could impede the rapid growth of the internet [7]. Another popular idea is to have more than one external links connected to a customer, which is termed as multi-homing. To implement a multi-homing network, one does not have to change the underlying network routing scheme. Common solutions include multiple-entry Domain Naming Services (DNS) and Network Address Translation (NAT). So each machine will have multiple network addresses, namely IP addresses in the DNS or the NAT table. As a result, a multi-homed machine can be reached from one of its network addresses [6].

Studies [1, 2, 3, 8] have demonstrated that multi-homing approach is more cost effective and is easier to implement than overlay networks. And performance gaining from multi-homing network is comparable to overlay networks. On the network security side, the conventional Viral Private Network (VPN) is limited to point-to-point structure [9], which makes it very slow to switch over to another network in case of single node or single link failure.

While recent studies are focused only either on reliability or on security, none of them have put both schemes together. Therefore, we are trying to study how to integrate VPNs with multi-homing networks.

The paper is organized as follows:

- Introduction to virtual private network technology
- Background information of Multi-homing network
- Problems when supporting VPN in a multi-homed network
- Descriptions of proposed multi-connection solution
- Conclusions and future works

Introduction to Virtual Private Network (VPN) Technology

A Virtual Private Network (VPN) is a common network mechanism to provide a secure end-to-end network connection. The idea is to first negotiate and setup a network tunnel between the two communication nodes. Usually a VPN server also connects to a RADIUS server to allow only authorized users have the privilege to establish such tunnels. The data will then be encrypted before it is transmitted over the network and will then be decrypted on the receiver side. Compare to dedicated private leased lines, VPNs are much cheaper and are ideal to many companies.

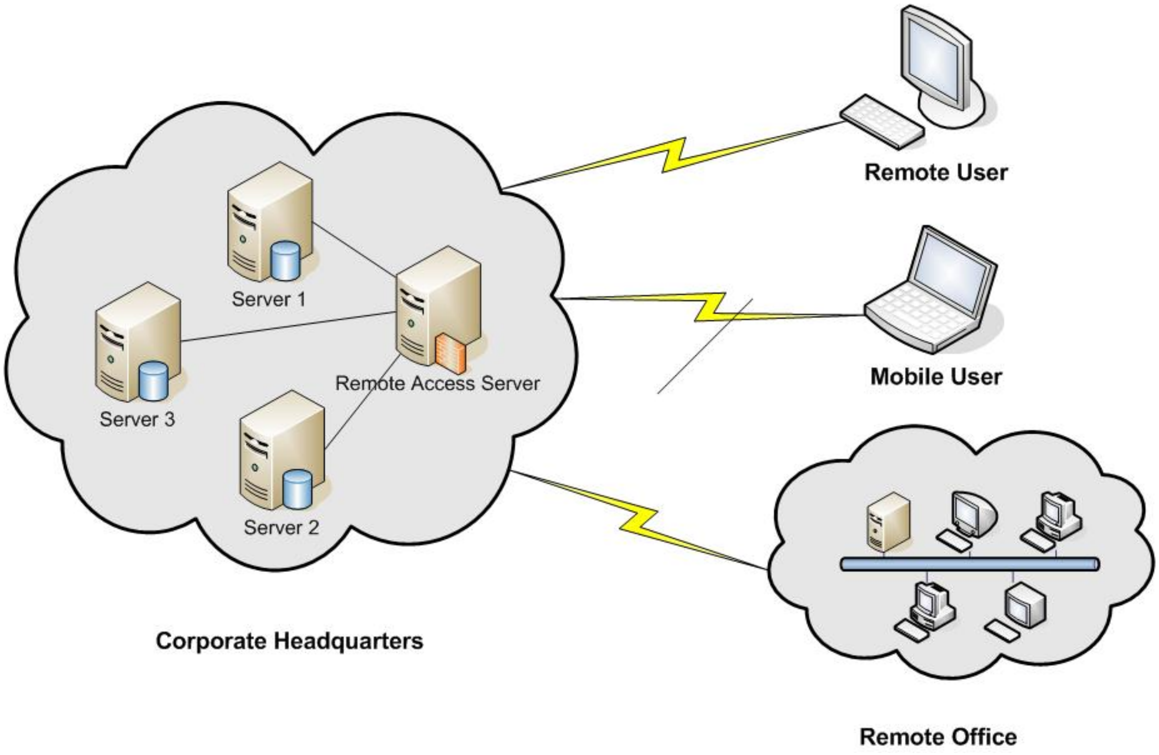

Figure 1: A network with out the support of VPN

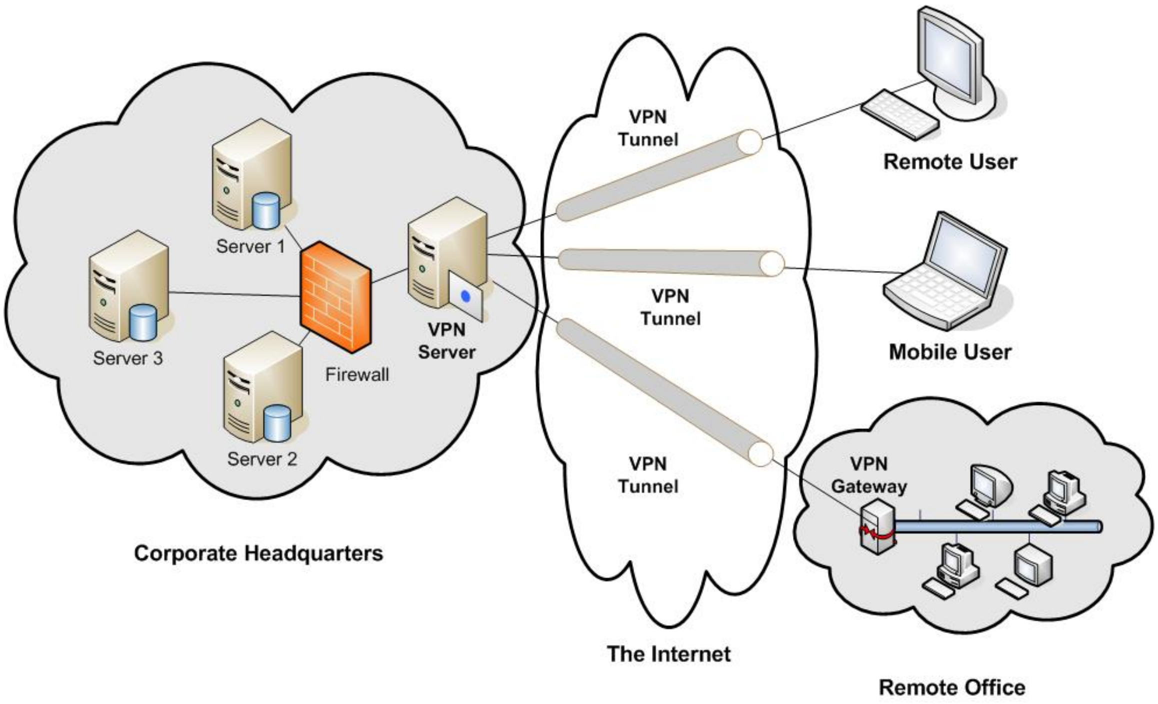

Figure 2: A Network with Virtual Private Networks

Consider a corporate network without the VPN support. Shown in Figure 1, a remote user needs to dial up to the corporate Remote Access Server (RAS) to get access to the servers. This is could be very pricy if users are trying to make a data connection through long distance calls. Similar problems also exist for mobile users and users at remote offices.

With VPN, however, users do not need to dial up directly to the corporate headquarter. Illustrated in Figure 2, a remote user can simply dial into the local access server and relying on the local ISPs to package the data and route the data through a 'tunnel' to the remote servers. Of course, you will need to pay for the tunnel services offered by the ISPs, which are normally less than half of what you pay for leased lines or long distance phone calls. Three major tunneling protocols are supported in the internet, as described below.

- IP Security (IPSec). It is developed by Internet Engineering Task Force (IETF) and operates at network layer. It can be implemented independent of application layer.
- Point-to-point Tunneling protocol (PPTP). This is the protocol developed by Microsoft, 3Come and Ascent Communications. It works at Data Link layer and is preferred for Microsoft Windows based network traffic.
- Layer 2 Tunneling Protocol (L2TP). It is the VPN implementation from Cisco. It combines Cisco's previous proposed Layer 2 Forwarding with PPTP. It offers more flexibility than PPTP, but need supports from the underlying network devices, such as routers and switches.

The advantages of VPN include reduced cost, effective use of bandwidth, enhanced scalability, and enhanced connectivity. It also has better security than conventional internet because it often encrypts the data during data transmission. The drawback of VPN is also obvious: it is highly dependent on the internet and is lacking interoperability of devices and protocols.

The background information of Multi-homed networks

Multi-homed network refers to the network with more than one link external to local network. For example, you can have two or more links to the same ISP to get link redundancy. In case one link fails, you still have another link operating. Or, to make it more reliable, you can sign contracts from different ISPs for the external links. The benefits of multi-homing include more bandwidth, high availability and cost effectiveness.

Several solutions have been proposed to support multi-homing technique. From end user perspective, it can be categorized as single IP solutions and multiple IPs solutions. The former is more user-friendly: a client machine does not need to worry about which path to take because the same IP addresses are registered at both paths of the egress router. An example of the single IP solution is shown in Figure 3.

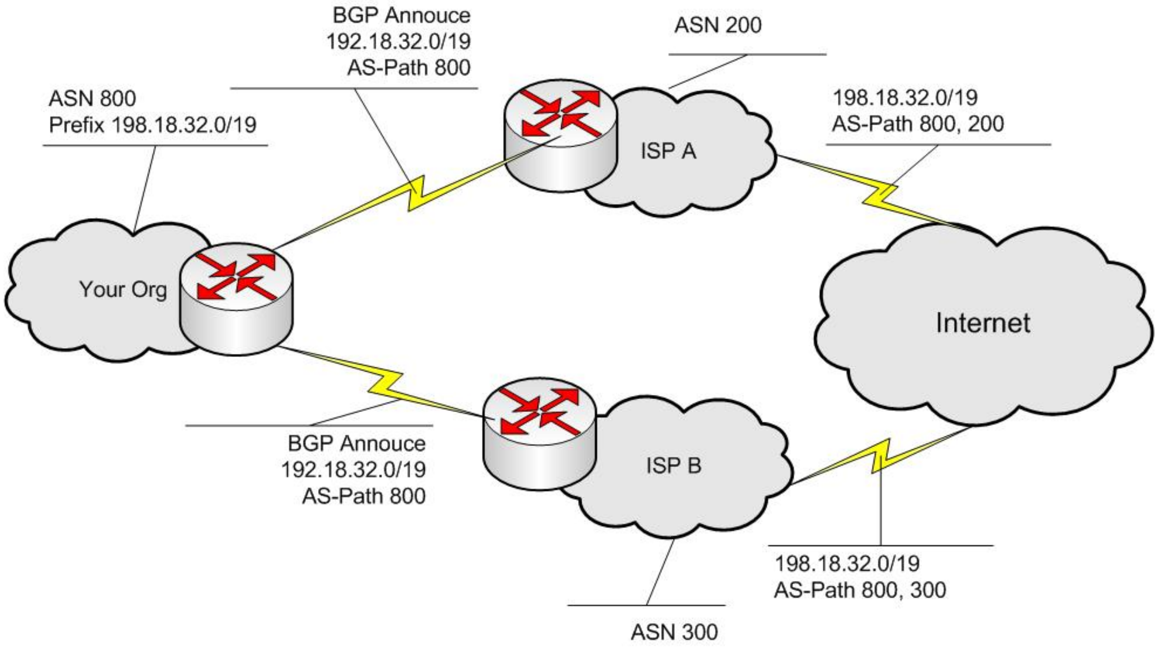

Figure 3: Multi-homing solution with Border Gateway Protocol offers same IP

In the example, an organization registered a network 198.18.32.0/19 at both ISP A and ISP B. Assuming Border Gateway Protocol (BGP) is supported on the routers of the organization and the access routers of ISP A and ISP B, the network announcement can be passed to the access routers of both ISPs. Then both access routers advertise the path to the internet. When the internet users try to reach the servers (e.g. web servers) of the organization, they see both paths are available from either ISP A or ISP B. Therefore, if one ISP has service difficulty, the network can still be accessed by the outside world from another ISP.

Another category of solutions offer different IPs with each from one unique ISP. For example, if a web server registered a domain name www.mywebserver.com. You can lease an IP of 143.120.18.25 from ISP A and an IP of 63.68.123.17 from ISP B. Bother IPs are then assigned to this very web server. The outside users can access the server through either 143.120.18.25 or 63.68.123.17. The difficulty of this solution is however, user has to know what IPs to use. Fortunately, this can be solved by multiple DNS entries.

And to map different IPs to the same network, Network Address Translation (NAT) is usually required. This is illustrated in Figure 4. In the example, a corporate network has two public networks address 63.136.32.16/29 and 147.226.16.8/29. Internally, the private network uses 10.1.0.0/16 network. A NAT device is sitting between the public and private interface to translate and map the internal IPs to external IPs, and vice versa.

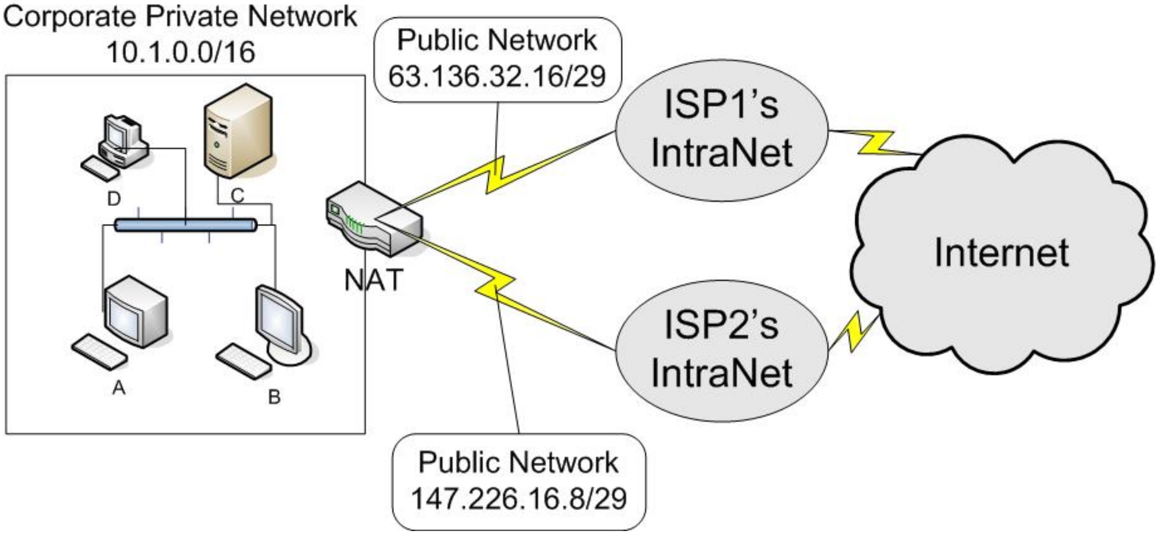

Figure 4: Multi-homing with NAT support

Supporting VPNs over Multi-homed Networks

As described in previous two sections, VPN technique offers cost-effective secure communication and multi-homing network provides an economy solution to enhance the network availability. Naturally, if there is a scheme that can incorporate both techniques, the resulting network infrastructure will be very attractive and enjoys the following features:

- Cost effective: Both VPN and Multi-homing are cost effective techniques.
- More secure: A verity of security features are available in the VPN world.
- More reliable: Because the network redundancy is provided by different ISPs.
- Better performance: More bandwidth is available, load-balancing is possible and Quality of Service (QoS) is possible.

Make no assumptions here, however, as the two will not go together automatically. The first issue we have identified is that the connectivity problem exists when incorporating IPSec with NAT. The problem occurs because both IP addresses and port numbers are encrypted in IPSec. When the encrypted IP and port number passed to a NAT device, it will not be able to decrypt the address and is therefore not capable of translating the IP address between internal and external networks.

In addition, auto fail over is a major issue. When one network path is not available, the ideal solution should re-route all the existing traffic through another links. However, in VPN technology, the tunnels are normally created end-to-end. When failure occurs, users need to re-setup a VPN connection to the live IP address.

Last but certainly not least, load balancing VPN traffic over Multi-homed network is challenging. Extra factor needs to be addressed since VPN requires more computing power especially when encryption and decryption are involved.

Multiple VPN tunnels solution

To address the auto fail over issue, there is a simple but not so elegant solution: educating the end users about the multi-homing. They do not necessarily have to understand the technology but the end users should realize when the VPN fails, they can create another VPN by using another IP address. The risk is, however, all the on-going transactions will be discarded. If end users have network translation that last for hours, this approach is not helpful.

If we aim at an elegant solution, we can create an “any-cast” type of VPN tunnel, which will send data through a tree like tunnel to any on of the public IP addresses of the VPN server. The difficulty is that any-cast is not even well supported on native IP networks and the behavior of tree-like tunnel is not well understood.

The solution we propose is to create multiple VPN tunnels simultaneously. The role of those tunnels can be classified as either primary/back up or a means for load balancing.

An illustration of our approach is presented in Figure 5.

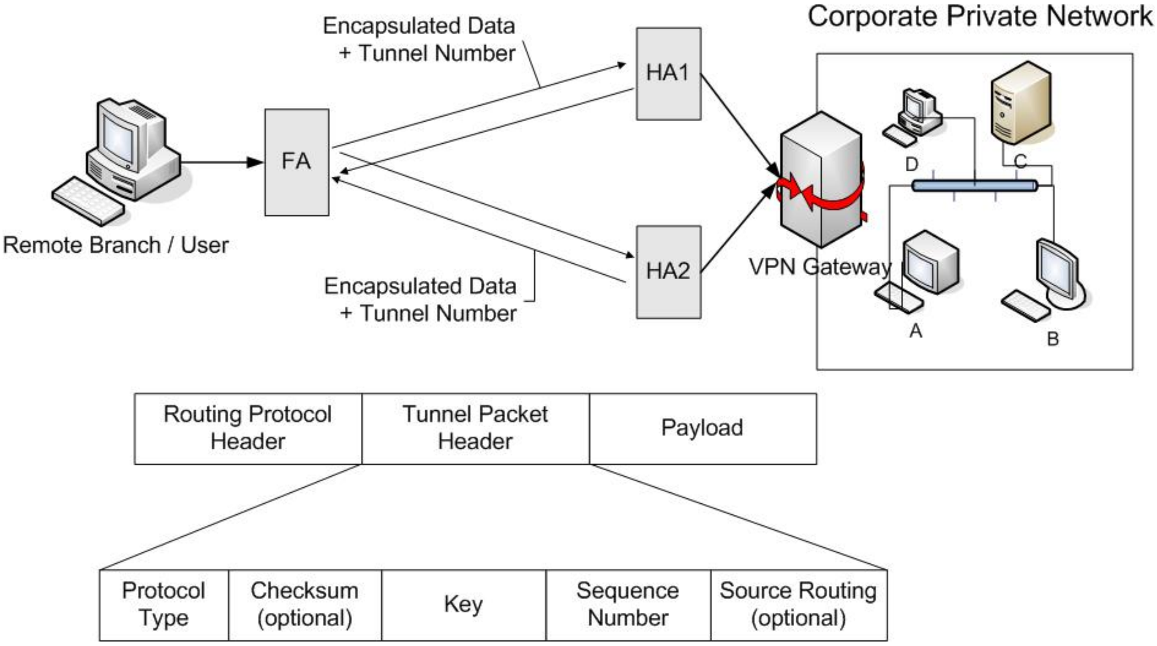

Figure 5: Multiple VPN tunnels illustration

In the preceding Figure, the home agent has two different public IP addresses, HA1 and HA2. When the end user tries to access the corporate private network from FA, two separate tunnels are created. Both tunnels insert tunnel packet header into the original IP packets, add tunnel number and encrypt the payload.

By default, multiple VPN tunnels do not provide auto fail over because home agents have different IPs. This can be solved by modifying the foreign agent's routing table when one path is congested or down. In our simulations, we modify the 'cost' entry of the routing table at the client side. The FA periodically evaluates the average of Round Trip Time (RTT) and changes the 'cost' of each path to reflect the network conditions. The fast link is therefore has precedence over the slow link and live link will always be chosen over the dead link. This solution not only solves the auto fail over issue, but also helps to perform load balancing and to provide QoS to the network application.

In terms of cost and performance, the proposed approach wastes no bandwidth. The only resource consumed is the maximum number of VPN connections a corporate network can support will drop by half. But it is a price you have to pay if the network availability or the fault tolerance is a concern. Auto fail over can be achieved from above analysis and the devices to implement the proposal are all off-the-shelf hardware requires no special computing powers. Of course, the cost of managing multiple VPN tunnels is slightly higher than maintaining a single VPN tunnel.

Conclusions and Future work

In this paper, we proposed a scheme to incorporate VPNs and multi-homed networks. The incorporation of the two techniques is challenging and beneficial. A number of issues are discussed in this context and a multiple-tunnel solution is proposed to support VPN on a multi-homed network. The proposed solution can be implemented either at networking layer or at application layer by changing the cost of the routing table on the fly. The resulting routing paths can provide auto fail over features and are helpful to balance the network traffic.

In the future, we would like to further test the performance of multi-tunnel solution on different platforms and a variety of networks to gain deeper understanding as how the scheme impacts the network and the end users. We would also like to propose a load balancing scheme that is more appropriate for VPN connections. In addition, a general solution to support NAT over IPSec will be examined. Lastly, point-to-multipoint tunnels will be investigated and constructed to further enhance the fault tolerance and to provide QoS to the end users.

References

[1] A. Akella, J.P., A. Shaikh, S. Seshan, and B. Maggs. "*A comparison of overlay routing and multihoming route control*". in proceedings of ACM SIGCOMM. 2004. Portland, Oregon.

[2] Armando L. Caro, J., Janardan R. Lyengar, Paul D. Amer, Gerard J. heinz, and Randall R. Stewart, "*Using SCTP Multihoming for Fault Tolerance and Load Balancing*", ACM SIGCOMM Computer Communication Review, 2002. 32(3): p. 23-32.

[3] David K. Goldenberg, L.Q., Haiyong Xie, Yang Richard Yang, Yin Zhang. "*Optimizing Cost and Performance for Multihoming*", in Proceedings of the 2004 Conference on Applications, technologies, architectures, and protocols for computer communications.

[4] Jahanian, J.H.a.F. "*Impact of Path Diversity on Multi-homed and Overlay Networks*", in International Conference on Dependable Systems and Netoworks. 2004. Florence, Itly: IEEE Computer Society.

[5] Ningning Hu, L.L., Zhuoqing Morley Mao, Peter Steenkiste, and Ja Wang. "*Locating Internet Bottlenecks: Algorithms, Measurements, and Implications*", in ACM Conference on Applications, technologies, architectures and protocols for computer communications. 2004. Portland, Oregen, USA: ACM Press.

[6] Ramakrishna Gummadi, R.G. "*Practical Routing-Layer Support for Scalable Multihoming*", in IEEE Infocom on Computer Communications. 2005.

[7] Ratual Mahajan, D.W., and Tom Anderson. "*Understanding BGP Misconfiguration*", in ACM SIGCOMM. 2002.

[8] Shu Tao, K.X., Ying Xu, Teng Fei, Lixin Gao, Roch Guerin, Jim Kurose, Don Towsley, and Zhi-Li Zhang, Exploring the Performance Benefits of End-to-End Path Switching. ACM SIGMETRICS Performance Evaluation Review, 2004. 32(1): p. 418-419.

[9] Z. Morley Mao, D.J., Oliver Spatsheck, Jacobus E. Van Der Merwe, and Jia Wang. "*Efficient and Robust Streaming Provisioning in VPNs*", in International World Wide Web Conference. 2003. Budapest, Hungary: ACM Press.

Biographies

BAIJIAN YANG is currently an Assistant Professor at the Department of Technology, Ball State University. He joined the Ball State in August 2003, after completing his Ph.D. in Computer Science at the Department of Computer Science and Engineering, Michigan State University. Dr. Yang's main teaching and research interests are in the field of computer networks, distributed systems and wireless communications.

TIANGUANG GAO is an Assistant Professor of Educational Technology at Ball State University. He joined the Ball State in August 2001, after completing his Ph.D. in Education at Purdue University. Dr. Gao teaches various courses in the field of educational technology, both online and on campus.