Postprint: Overlay Network Techniques for Inter-Domain Routing Optimization

Authors: Li Yanjun, 张国清

Date: 2017-03-09T00:00:00+00:00

Abstract

The distributed management of the Internet leads to an inherent defect in that its overall transmission quality cannot be guaranteed, and inter-domain routing optimization overlay networks have emerged as a prominent technology to address this issue. This paper reviews the current state of Internet routing, provides a refined classification of existing overlay network types, and analyzes the advantages and disadvantages of each category of technology. Building upon this foundation, it explores various issues confronting inter-domain routing optimization overlay network technology, such as network conflicts and traffic oscillations, and performs an in-depth quantitative comprehensive analysis of overlay network routing performance improvements. Finally, it summarizes the key technical challenges involved in implementing a practical, efficient, and scalable inter-domain routing optimization overlay network.

Full Text

Inter-Domain Routing Optimization Overlay Network Technology

Li Yanjun, Zhang Guoqing

Abstract

The distributed management of the Internet has led to inherent defects in guaranteeing overall transmission quality, making inter-domain routing optimization overlays a hot topic for addressing this issue. This paper reviews the current state of Internet routing, provides a refined classification of existing overlay network types, and analyzes the advantages and disadvantages of each category. On this basis, we discuss various problems facing inter-domain routing optimization overlay technologies, including network conflicts and traffic oscillations, and conduct an in-depth quantitative and comprehensive analysis of overlay routing

performance improvements. Finally, we summarize the key technical challenges in implementing a practical, efficient, and scalable inter-domain routing optimization overlay network.

1 Introduction

The development of the Internet, centered around TCP/IP technology, originated from military requirements for survivability under attack. It emphasizes the independence, distribution, and autonomy of individual subnets. After more than 20 years of commercial evolution, the Internet has established an unshakable position in the field of network communications due to its excellent uniformity and compatibility, with explosive growth and applications that even its designers never anticipated.

Today, the scale of the Internet continues to expand rapidly. Although data exchange and routing hardware and software technologies are constantly innovating and advancing, the characteristics of distribution and independence have remained unchanged since its inception. The vast Internet is controlled and managed by many different organizations and operators, and the conflicts of interest and differences among them have become serious constraints on Internet evolution. Numerous global failure incidents in recent years, such as the Taiwan submarine cable break and the Pakistan Telecommunication Authority' s blocking of YouTube, have revealed that the Internet, which appears increasingly mature and complete on the surface, remains extremely fragile.

In response, relevant Internet organizations and institutions have made many efforts to overcome this fragility and provide corresponding transmission quality guarantees for various services. As early as the 1990s, the IETF recognized these defects in the Internet and proposed various overall Quality of Service (QoS) architecture solutions such as DiffServ and InterServ [1][2]. However, implementing these solutions requires two basic conditions: first, all routers along the communication path must perform unified and coordinated service scheduling and buffer control; second, when end-to-end network traffic traverses multiple network domains maintained by different Internet Service Providers (ISPs), active cooperation and coordination from ISPs are required. For commercial reasons, however, ISPs only hope to obtain maximum market returns with minimum investment, and each ISP is only willing to provide service guarantees for transmission quality and survivability within its own network domain, unwilling to increase burdens for improving overall service quality.

Therefore, to date, almost no effective overall solution has been implemented on the existing Internet. On the other hand, a large number of emerging applications with strong real-time requirements, such as audio-video interaction, IPTV, and online gaming, have increasingly become mainstream network content, and their demand for overall QoS from IP networks is growing stronger. How to provide satisfactory overall data transmission performance for services spanning multiple network domains (i.e., inter-domain transmission) has become one of the most challenging focal issues in current network technology research.

Obviously, in the current vast network environment, research limited to a single layer of IP routing protocols can hardly solve the problem. New-generation network routing technology needs to seek an intermediate layer model mechanism that can provide end-to-end service capabilities across different networks, operators, and management domains while ensuring forward compatibility with existing IP protocols. In this environment, overlay networks built on top of the IP layer have emerged and gradually become the mainstream approach to solving these problems. Over the past decade, foreign research institutions have proposed numerous specific network structures and construction schemes, such as Detour, Resilient Overlay Network (RON), QRON, SON, NIRA (New Inter-Domain Routing Architecture), and others. However, due to various constraints, none of these schemes have achieved large-scale deployment, and network development once again hesitates at the crossroads.

This paper attempts to classify and discuss these inter-domain routing optimization technologies, summarize their advantages and disadvantages, and based on this, conduct an in-depth review and analysis of the feasibility and key challenges of overlay network technology in inter-domain routing optimization. The remainder of this paper is organized as follows:

Section 2 reviews the research status of existing routing optimization overlays and provides a complete classification of overlay types, summarizing the technical advantages and disadvantages of each type. Section 3 discusses various problems and challenges facing inter-domain routing optimization overlay technology. Section 4 analyzes the key technical issues in implementing an efficient and scalable inter-domain routing optimization overlay network. Section 5 presents conclusions and future outlook.

2 Overlay Network Research Status and Classification

The concept of overlay networks can be traced back to the early Internet. In a sense, the Internet itself can be viewed as a huge overlay network, using the IP layer as the overlay layer and leveraging IP protocols to provide "best-effort"data transmission services across heterogeneous underlying networks such as Ethernet and Token Ring. The rise of overlay networks stems from the fact that the existing Internet' s inter-domain routing capabilities have become increasingly unable to meet the needs of various carried services.

2.1 Current State of Internet Inter-Domain Routing

As early as 1995, V. Paxson et al. began statistical studies on the emerging Internet' s inter-domain routing performance [3], finding that the Internet had approximately 3.3% severe routing failure rate, with over 20% of path failures unable to be repaired within 10 minutes. C. Labovitz et al.' s statistics from 1997 to 2000 showed that 10% of routers in the existing Internet had reliability below 95%, 65% had reliability below 99.9%, and failure recovery times often reached 15 minutes, with over 40% of routing failures requiring more than 30 minutes to recover [4]. Chandra et al. discovered through active probing that 5% of failures in the Internet could even persist for over 2.75 hours [52].

For historical reasons, the existing Internet is actually a vast network formed by interconnecting many relatively independent Autonomous Systems (AS) through the inter-domain routing protocol BGP (Border Gateway Protocol). Each autonomous system is generally controlled by an independent entity, mostly commercial organizations such as ISPs. The BGP protocol is the current border gateway routing protocol used between Internet autonomous systems. As the core component of the Internet, it is the basic mechanism and interconnection bond for passing network information between all autonomous systems, as well as the primary means for ISPs to implement policy control, playing a crucial role in Internet evolution. However, since its creation in 1989 and revision to the fourth version in 1995, the BGP protocol has hardly undergone any major changes.

For scalability considerations, the BGP protocol only uses simple hop count as a metric, which results in the so-called best path not being the optimal routing performance path [5]. To avoid "route flapping" caused by frequent path updates, BGP does not perform real-time path performance probing and uses single path selection. This prevents business traffic from avoiding congestion points when links become congested, and after failures occur, it exhibits severe slow convergence characteristics [6].

These measurement results and research conclusions have fully demonstrated that the current end-to-end Internet inter-domain routing mechanism neither guarantees optimal transmission quality between end hosts nor enables rapid failure recovery or response. To compensate for BGP' s defects, researchers have proposed various extension schemes for the BGP protocol [8-10]. However, these ideas and suggestions all require large-scale additional upgrade deployments for implementation, lacking transitional feasibility.

2.2 Overlay Network Application Status and Classification

2.2.1 Overlay Network Application Status To overcome the inherent defects of Internet inter-domain routing and meet the security, stability, and reliability requirements of emerging services, people have constructed various specialized overlay networks for specific service types, such as Content Distribution Networks (CDN) for publishing and storing streaming media data [11], end system multicast (ESM) technology [12] and Overcast systems [13] for application-layer multicast, Peer-to-Peer (P2P) networks for file sharing [14], the Plantlab simulation experimental network [15], and various combined service networks [16].

These networks are virtual networks composed of a series of service nodes distributed within various autonomous systems across the Internet and logical links connecting them. Through the forwarding of service nodes, user terminals can more effectively utilize Internet resources to provide performance requirements that existing Internet routing cannot satisfy for corresponding services. The greatest advantage of this overlay network implementation approach is that it can be flexibly and conveniently deployed on top of the basic network without requiring changes to the existing network architecture.

In recent years, inspired by research on these application overlays, overlay networks supporting QoS routing have rapidly become a hot technology for solving problems such as slow fault convergence and poor stability in Internet cross-domain transmission. As an intermediate layer model mechanism, routing overlays can actively select efficient paths to route data packets through node cooperation, thereby improving end-to-end network transmission quality, meeting user application performance requirements, and providing more reliable and fault-tolerant services [17][18][19]. In typical routing overlay technologies such as Resilient Overlay Networks and Detour [20], researchers have verified through building corresponding experimental networks the tremendous advantages of using routing overlays for cross-domain transmission in terms of rapid response, fault recovery, and QoS guarantees compared to the BGP protocol.

2.2.2 Classification of Inter-Domain Routing Optimization Overlays

Based on different overlay implementation technologies, these cross-domain QoS routing support overlay networks can be divided into two major categories: relay-type overlays and Virtual Network (VN)-type overlays. According to the location of networking nodes and system implementation layers, they can be further subdivided into network-layer overlays and application-layer overlays.

(1) Relay-Type Routing Optimization Overlays

The principle of relay-type overlays is that the sender transmits packets to a relay node, which then forwards them to the receiver, thereby forming a multi-hop overlay path. Since two-hop overlay paths provide performance, reliability, and path diversity similar to multi-hop overlay paths [21], most overlay research adopts two-hop overlay paths forwarded through a single relay node for simplification and ease of implementation.

Resilient Overlay Networks (RON)

Resilient Overlay Networks are a typical representative of relay-type overlays [17]. In terms of system implementation layer, it is an application-layer overlay built on top of the existing Internet routing layer. Applications built on it can utilize the RON routing library established on the underlying network to make path selections on the existing Internet. The architecture of RON is shown in [Figure 1: see original paper]. It adopts a two-hop relay application-layer routing approach. System nodes monitor path functions and performance between each other, regularly checking network connectivity between themselves and other nodes. The application layer can evaluate multiple forwarding paths generated by selecting different transit nodes based on comprehensive multi-dimensional metrics such as delay, packet loss rate, and throughput, and then select forwarding paths that meet service QoS requirements. At the same time, RON's application-layer routing approach can also utilize hidden paths that the BGP protocol cannot discover or utilize during cross-domain transmission for rapid routing fault recovery, improving packet transmission reliability. Although RON achieves good transmission optimization effects between point-to-point connections and can react quickly in small-scale network environments, it has many shortcomings:

First, RON nodes must send probe packets at fixed intervals in a fully meshed connection manner, creating a high-intensity probing load. This small-interval, frequent probing brings huge overhead to the network, greatly limiting RON' s scalability.

Second, RON' s virtual network uses static management, with forwarding nodes limited to a few server nodes, making the structure inflexible and unable to fully and effectively utilize hidden paths and private paths.

Finally, RON' s path selection is only based on performance comparison among neighbor nodes currently probed via flooding, without reference to the underlying topology. Therefore, the correlation between overlay paths and underlying default IP paths is high, making it powerless against certain performance degradations and path failures that occur on the Internet.

QRON (QoS-Aware Routing in Overlay Networks)

To address RON' s scalability and performance optimization issues for specific services, Z. Li et al. proposed the QRON architecture [22], attempting to use a generic Overlay Service Network (OSN) to meet the needs of various application-layer services.

QRON utilizes Overlay Brokers (OBs) to establish the overlay service network. These OBs are deployed by overlay service providers in various autonomous systems across the Internet, cooperating with each other to form an overlay service network that provides services to upper-layer applications, such as resource allocation and negotiation, overlay relay routing, and topology discovery. QRON focuses on designing a QoS-aware routing protocol, with the main purpose of finding overlay paths that meet QoS requirements for upper-layer QoS-sensitive overlay applications, while simultaneously performing overlay traffic balancing. Since OBs in QRON adopt a hierarchical organization and are clustered based on "network distance," QRON has relatively good scalability.

Compared with RON, QRON is a relay overlay with virtual network properties. Its hierarchical deployment offers better scalability; its routing algorithm based on path performance makes path selection more efficient; and its design of a generic QoS-aware routing protocol provides better universality. However, QRON is a network-layer overlay, and all its nodes require self-deployment, making the cost very high. Additionally, when evaluating path performance, QRON only measures packet loss, delay, and available bandwidth, and the measurement methods are not highly accurate. This routing algorithm designed with inaccurately measurable available bandwidth greatly limits its universality.

Spines Overlay Network

Spines is a generic overlay network system developed by the Distributed Network Laboratory at Johns Hopkins University [25][26], with open source code that can be used for testing and developing overlay network protocols. Spines provides a two-level hierarchical structure. User programs connect to the nearest overlay node. This overlay node sends or forwards data to the destination address through the overlay network. Spines runs at the application layer, requiring no changes to infrastructure and core access, thus having good deployability like RON. The benefit of its two-level structure is that it can limit the scale of the overlay network to reduce the amount of data exchanged for flow control information. Overlay nodes serve as both servers for connecting various applications and as forwarding routers for transmitting data packets between nodes. Due to this hierarchical structure of Spines, applications can be located either at overlay nodes or on other machines connected to the nodes. Obviously, Spines has made great progress in scalability compared with RON. Through the Spines system, services can achieve low-overhead reliable transmission, improve real-time transmission characteristics such as delay, packet loss rate, and throughput, and support multimedia services with high QoS requirements, such as VoIP and video conferencing.

However, the Spines network also has its own disadvantages: on the one hand, Spines only supports fixed-node topology structures, and nodes cannot dynamically join or leave the Spines network. Once a node or link fails, nodes may need to take very long detours to transmit information, or the entire network may collapse. On the other hand, the topology structure established by Spines is often not optimal, and it has no understanding of the actual topology of the underlying network, resulting in relatively poor transmission flexibility and efficiency.

(2) Virtual Network-Type Routing Optimization Overlays

Like relay-type overlays, virtual network-type overlays also place nodes at key network locations. However, nodes in virtual network-type overlays differ from relay-type nodes in that they also have functions such as monitoring underlying network resource distribution, bandwidth utilization, link performance, and traffic load, and can perform dynamic resource allocation based on monitoring results to achieve overall network performance optimization. Virtual network-type overlays can be understood as “virtual overlay private networks,” within which end-user traffic can receive transmission optimization.

Detour Routing

The Detour routing system proposed by S. Savage et al. in 1999 can be considered an early virtual network-type routing overlay. It consists of a set of distributed edge routing nodes connected through tunneling technology. Traffic data entering Detour’ s ingress node is re-encapsulated into new IP packets, transmitted along the Internet to Detour’ s egress node, and finally delivered to the user end (as shown in [Figure 3: see original paper]).

Detour attempts to study performance-sensitive new routing algorithms on a large-scale network to improve Internet network performance. It mainly includes two functions: virtual network management and node routing control. User nodes enter the Detour virtual network by changing their default router to a nearby edge Detour node, which then transmits traffic to the destination node.

Compared with Internet routing, the Detour system performs router selection adjustments on a minute-level time scale. By using techniques such as detour routing, single-flow multi-path transmission, and packet classification, it can improve end-to-end transmission performance to some extent. However, its overhead and limitations are also obvious: first, Detour’ s virtual network is deployed above underlying routing nodes and runs independent algorithms. To form a unified large-scale network applicable across the Internet, the deployment cost would be extremely high. Second, even if the default Internet routing is more efficient, all user traffic using Detour will still be transmitted through the Detour network. This increases Detour’ s load, causing it to face the same old problems of scalability and efficiency contradictions as the original Internet. Additionally, Detour’ s network structure lacks flexibility and cannot fully and effectively utilize the large number of hidden paths on the Internet.

OverQoS

If the Detour network still has characteristics of relay-type overlays, then OverQoS [23] is a very typical virtual overlay network. In the OverQoS architecture, service providers purchase network access rights from traditional ISPs and then deploy OverQoS nodes in different routing domains to form an overlay network. In traditional QoS mechanisms, IP routers are responsible for controlling packet buffers and output link bandwidth, while OverQoS nodes do not consider underlying performance. It introduces the concept of Controlled-Loss Virtual Link (CLVL).

By controlling the aggregation speed of flow bundles, the packet loss rate can be maintained within a small range. Therefore, regardless of how the network environment changes, OverQoS can obtain a lower-bound QoS service, thereby achieving statistical guarantees on bandwidth and packet loss rate metrics.

The OverQoS network belongs to network-layer overlays, so it has the same deployment cost issues as Detour. At the same time, OverQoS’ s optimization of stream transmission mainly involves smoothing packet loss and specifying packet priorities to statistically guarantee bandwidth, without considering other transmission performance metrics. This limits the effectiveness of application transmission optimization.

SON (Service Overlay Networks)

SON [24] is an overlay network for solving end-to-end QoS and constructing and deploying QoS-sensitive services. It provides a logical end-to-end service transmission platform by purchasing network bandwidth with QoS guarantees from various network domains. SON provides value-added network services to users through service level agreements and fee collection. SON’ s structural concept is also closest to the virtual network model.

3 Inter-Domain Routing Optimization Overlays

Although existing research shows that routing overlays can improve service transmission quality to some extent by utilizing redundant Internet paths in various ways, truly deployable inter-domain routing optimization overlays still face many practical problems that need to be solved. First, in terms of implementation difficulty, to solve existing Internet inter-domain routing problems, overlay approaches need to be widely compared with numerous alternative technical solutions such as rebuilding protocol layers or using multi-homing to establish their advantages. Second, from a management perspective, as a supplement to the Internet, routing overlays will face the same scalability and manageability issues as the Internet once deployed on a large scale. It is necessary to carefully consider whether their traffic optimization functions will conflict with ISPs’ IP-layer traffic management strategies and with other overlays, causing network oscillations. Third, the overlay scheme’ s own ability to improve network traffic performance and its overhead need further quantitative comparison. Finally, there is disagreement over whether overlay technology should be implemented by upgrading routers or deploying terminal servers.

3.1 Comparison of Inter-Domain Routing Optimization Overlays with Other Optimization Schemes

Regarding solutions to current Internet transmission service problems, there has been significant disagreement: one type of solution advocates a complete overhaul, such as GENI and FIRE; another advocates full compatibility with existing network architectures, such as overlays and multi-homing; and a third type takes a compromise, believing that some local architectural innovations can be made.

3.1.1 Overlay Networks vs. GENI and FIRE GENI is a next-generation experimental research network project promoted by the U.S. National Science Foundation (NSF) [37]. Unlike overlays that meet compatibility and transition requirements, GENI considers “global network needs after 15 years” and how to build networks unconstrained from scratch (Clean-slate). Its research directions are not limited to infrastructure but also include principles, protocols, architecture, and design solutions. Currently, schools and research institutes participating in GENI include almost all top-tier U.S. institutions: Stanford, MIT, Princeton, and even the Department of Defense. Actively participating companies include Cisco, Fujitsu, HP, Microsoft Research, NEC, and many other well-known international enterprises. GENI’ s core conceptual requirements for architecture are: programmability, virtualization, and resource sharing.

The EU has a similar project in the direction of network architecture innovation design called FIRE (Future Internet Research and Experimentation) [38]. FIRE' s principle is to promote experiment-based research, combining forward-looking academic research with industrial testing and experimentation. Similarly, FIRE is also committed to creating large-scale, dynamic, and sustainable experimental facilities in Europe, a process that will link and unite various small-scale testbeds currently in existence.

However, whether GENI or FIRE, these "clean-slate" national research efforts are still in their early stages—it may be difficult to expect meaningful results within 10 to 15 years. Compared with overlay technology, both GENI and FIRE represent a completely new revolution, but from a practical standpoint, evolution may be a more realistic choice. Because today' s Internet far exceeds what academia alone can control, starting from scratch faces far more difficulties today than the Internet' s creators faced 40 years ago.

3.1.2 Overlay Networks vs. Multi-Homing Compared with overlay deployment, multi-homing is a simpler performance optimization method. Large enterprises or small ISPs typically connect to the Internet in this way to obtain more reliable and stable routing performance. However, using multi-homing to access the Internet requires paying ISPs for redundant link usage fees, with optimization effects directly proportional to cost. D.K. Goldenberg et al. established a cost model combined with performance metrics such as delay to design an intelligent multi-homing algorithm [39]. Through this algorithm, network performance can be improved by about 18% at the cost of adding one ISP connection on top of single-homing.

A. Akella conducted a detailed comparison of the optimization capabilities of multi-homing and overlays relative to default routing in terms of Round-Trip Time (RTT) and throughput [40]. The tests showed that with a single overlay and single multi-homing, overlays have better optimization capability, with delay metrics averaging 33% better and throughput 15% better. With K multi-homing (K>3), multi-homing's network performance is slightly better than overlay technology compared to a single overlay. If overlay technology is used on top of multi-homing (K>3), performance improvement is very limited (about 10% delay improvement and 5% throughput improvement). At the same time, for ISPs to use overlay technology, they also need to sign complex traffic forwarding agreements with each other. Therefore, Akella believes that ISPs can completely improve network performance through intelligent routing mechanisms based on multi-homing without needing to specifically build overlay networks.

However, although multi-homing implementation is simple, multiple connections cause client costs to multiply, while overlays can provide a large number of discrete paths after scaling up, further improving optimization cost-effectiveness. Therefore, Yong Zhu et al. proposed in [41] that the advantages of both can be combined to form a MON (Multihomed Overlay Network) and conducted in-depth discussions on MON' s deployment location, design methods, and operational overhead. Their theoretical analysis results show that using the Overlay Service Provider (OSP) operation model can not only improve comprehensive network performance but also reduce operational costs and increase profit margins compared to traditional ISPs. However, since MON design is a complex NP problem, obtaining optimal optimization results requires collecting and analyzing global information on users, traffic, and performance, which to some extent limits its application. Therefore, H. Okada et al. further studied cooperative routing methods for multiple coexisting overlays and proposed a relatively simple cooperation protocol, but its effectiveness still needs further verification in real network environments [42].

3.1.3 Overlay Networks vs. Other New Architectures Inspired by the active routing characteristics of overlay technology and multi-homing, researchers have proposed some compromise network transformation schemes, with a typical representative being the NIRA (New Internet Routing Architecture) proposed by X.W. Yang et al. [43].

Similar to overlay technology, NIRA is a user-oriented active routing model. By designing a new Internet inter-domain routing architecture, NIRA enables end users or applications to select the series of ISPs that their data packets traverse. By giving end users more routing control at the domain level, on the one hand, it can force ISPs to compete with each other to improve service levels, and on the other hand, it can bring technical benefits and improve routing. NIRA uses TIPP (Topology Information Propagation) for topology information dissemination to achieve route discovery; adopts hierarchical addressing for efficient route description; and uses NRLS (Name-to-Route Lookup Service) for name-to-route lookup conversion when users initiate communication. NIRA has developed a complete routing system based on these key technologies that can support user route selection.

Although NIRA has drawn on research experience from multi-homing and overlay technology in routing improvement [44], avoiding their respective shortcomings, and its hierarchical addressing approach can greatly reduce routing state recorded in routing tables, solving scalability issues.

However, NIRA requires deployment on backbone networks, which is very costly. Additionally, its end users need to perform complex configurations, users have difficulty predicting service quality, and there are risks of exposing ISPs' routing business secrets. These factors all limit NIRA' s application and scalability.

3.2 Network Conflicts in Cross-Domain Routing Optimization Overlays

Cross-domain routing optimization overlays are typically composed of standby nodes distributed in different network domains, performing QoS control through application-layer routing or virtual routing networks. The idea is to transfer routing decisions to the application layer, which selects corresponding paths based on its own requirements for end-to-end delay, packet loss rate, or throughput, thereby breaking the IP layer’ s monopoly on routing.

However, allowing both the service layer and IP layer to simultaneously control traffic routing may trigger conflicts and contradictions between various independent traffic management and optimization strategies [27]. For example, in the network structure shown in [Figure 5: see original paper], the backup path A→C→F→G for communication node pair A-G in overlay 1 overlaps with overlay 2 (C→F), while the backup path for C-K in overlay 2 also overlaps with overlay 3 (B→E). When path D-G fails or becomes congested, overlay 1 will shift its upper-layer traffic to the backup path (A→C→F→G), thereby increasing traffic load on link C-F. If this traffic migration reduces communication quality between C-K, overlay 2 will shift its traffic to its backup path (C→B→E→H→K). Similarly, since overlay 2’ s backup path contains link B-E that overlaps with the default path between communication node pair B-H in overlay 3, once this migration affects communication between B-H in overlay 3, overlay 3 will shift traffic to its backup path (B→A→D→F→H). This is just the traffic migration triggered by link performance changes in a simple topology with 3 overlays and 9 nodes. In cases with more overlays and larger topology scales, traffic fluctuations triggered by single events become more complex and may even cause large-scale network oscillations due to repeated path switching [28].

We can summarize the network conflicts caused by traffic optimization in overlay networks into two categories: horizontal conflicts and vertical conflicts. Horizontal conflicts refer to conflicts between traffic from one overlay and traffic from other overlays or background traffic, as shown in the figure above. Vertical conflicts refer to conflicts between overlay traffic and ISP traffic management strategies. ISPs use Traffic Engineering (TE) strategies to cope with dynamic changes in network performance (such as link failures, BGP failures, and traffic congestion). Traffic engineering generally performs traffic balancing through dynamic estimation of traffic matrices (TM) [30]. ISPs typically make two assumptions about traffic engineering: (1) traffic demand will not change significantly in a short time; (2) changing transmission paths within the domain will not change the traffic demand itself. Inter-domain routing overlays precisely violate these two assumptions.

3.2.1 Horizontal Conflicts Among Multiple Overlays Regarding the horizontal and vertical conflict problems generated by overlay optimization, numerous researchers have conducted in-depth theoretical analysis and experimental verification. Regarding horizontal conflicts, most researchers believe that coexistence competition among overlays based on greedy optimization strategies will reduce overall network performance [31][32][33]. On this basis, Wenjie Jiang et al. used a non-cooperative game model to further conduct theoretical analysis and simulation experiments on overall optimization strategies, overlay-based optimization strategies, and selfish optimization strategies in multiple overlays sharing underlying links. The results suggest that overlay-based optimization strategies have performance relatively close to the optimal solution (overall optimization strategy), while the selfish optimization strategies adopted by existing overlays perform worse than the aforementioned two strategies in terms of both overall network performance and individual service performance of single overlays [34].

However, L. Qiu's simulation results in a single autonomous system environment using the OSPF protocol [29] show that in terms of horizontal conflicts, overlays do not degrade performance as much as the aforementioned research conclusions suggest. In an Internet-like environment, when each independent overlay adopts selfish routing optimization methods, at game equilibrium, both average network delay and link utilization efficiency are relatively close to optimal results, but it leads to system overhead caused by frequent overload of some shortest paths. Even increasing the number of overlays has only a very slight impact on end-to-end average delay.

This indicates that after adopting existing overlays' selfish routing methods, the impact of horizontal conflicts on performance degradation is not significant. Qiu' s simulation results also show that when network-layer routing mechanisms are reasonably configured, different overlay routings can coexist well.

The fundamental reason for these divergent research results lies in differences in comparison objects and simulation environments. Researchers who believe that overlays' selfish routing methods degrade overall performance base their comparison algorithms on the premise of global information sharing and centralized control management, and basically do not consider the impact of network-layer protocols (such as OSPF, MPLS, etc.) on overlay optimization algorithms. When Qiu et al. measured the impact of horizontal conflicts on performance in multiple overlay optimizations, they used a simulation environment close to the real Internet, comparing delay and network throughput performance metrics with two methods: optimization control based on information sharing and default passive Internet routing.

Obviously, achieving large-scale global detailed information sharing or centralized routing in the current Internet environment will be greatly constrained in terms of feasibility. Relatively speaking, implementing an independent overlay system for specific service routing requirements appears more practical. Therefore, Ram Keralapura et al. conducted further in-depth research on resource synchronization competition phenomena in overlay coexistence [28]. The results show that even if overlays adopt different path performance probing methods, they may still lead to synchronized competition for optimal links. This competition not only affects the optimization performance of each overlay but also affects non-overlay traffic in the network. By adding random parameters to the probing algorithm and using exponential backoff resource competition algorithms, the probability of overlay synchronization and synchronization oscillations can be effectively suppressed.

3.2.2 Vertical Conflicts Between Overlays and ISP Networks Regarding whether overlay applications will conflict with ISP network traffic engineering and reduce network performance, researchers' conclusions are relatively consistent. Qiu' s research shows that after introducing dynamic traffic matrix Tt(s,d), due to overlays' proactive adjustment of their own traffic, the traffic matrix changes frequently, increasing the difficulty of underlying optimization for traffic engineering strategies. The conflicting results of their respective optimizations instead reduce overall network performance. At the same time, Qiu also found that when using the MPLS protocol for underlying traffic optimization, it has better coexistence capability with overlays compared to OSPF. The reason is that MPLS has richer means of adjusting routing traffic.

Y. Liu et al. built a simulation environment for an overlay system completely isolated from the underlying network and using independent optimization strategies [35], assuming that overlays and ISP traffic engineering perform same-frequency Nash games. On this basis, they found that if overlay systems frequently optimize their performance, it would instead greatly increase their overhead, ultimately leading to degraded optimization performance. This is because after overlays select paths, the dynamic traffic matrix will adjust its own routing to minimize overall network overhead, while updated traffic engineering strategies will increase overlay overhead. This interactive game will repeat until Nash equilibrium is reached. Liu' s simulation experiments also show that compared with OSPF, traffic engineering using the MPLS protocol can achieve better optimization effects. At the same time, when overlay traffic dominates, its routing decisions have a very obvious impact on traffic engineering overhead.

Although these research results indicate that ISPs using MPLS-based traffic engineering strategies could theoretically avoid vertical conflicts when overlays perform traffic optimization, many challenging problems remain in practical application, such as obtaining dynamic global traffic matrix information and large-scale linear programming implementation. Therefore, ISPs'attitude clearly tends to restrict the development of overlay applications to reduce network overhead and prevent overall performance degradation.

However, J.H. Wang et al., from a business model game perspective [36], discovered through modeling and derivation of intra-domain and inter-domain traffic that large ISPs will fall into a dilemma at Nash equilibrium: when the traffic model migrates from traditional Web services to overlay methods, it not only leads to network traffic instability but also allows small ISPs to enjoy free-ride benefits. Simply reducing or limiting overlay end-user traffic will damage ISP business, affect access pricing power, and lead to customer loss; reducing private link usage will affect other services and network expansion, which is not conducive to their own and the Internet' s development.

In summary, overlays have become an irreversible technical trend. Their emergence strengthens the loose coupling of the Internet based on the BGP protocol, different autonomous systems will affect each other, traffic matrices become more dynamic and complex, and traditional traffic balancing strategies will be broken. ISPs must fully understand overlay networks and, based on this, adopt corresponding countermeasures, conduct characteristic pattern statistics, and provide good interactivity and information sharing mechanisms. Overlays should also understand ISPs' routing strategies and underlying information as much as possible and adopt corresponding mechanisms such as random path probing to avoid synchronization oscillations and prevent adverse effects on other overlay traffic and network background traffic (non-overlay traffic).

3.3 Practical Foundation and Improvement Capability Analysis of Inter-Domain Routing Optimization Overlays

3.3.1 Practical Foundation of Inter-Domain Routing Optimization Overlay Technology The important means and strategy for overlay networks to perform network optimization is to find more reliable backup paths for applications, allowing users to switch application transmission paths to these backup paths when necessary to ensure transmission QoS. The premise of this method is that the underlying network must have multiple redundant paths between communication node pairs, especially multiple independent physical paths. The current network can indeed provide such redundancy.

[Figure 6: see original paper] shows a cumulative distribution of independent shortest paths between all node pairs measured on 1235 edges among 42 nodes on RouteView [45]. From the figure, it can be seen that 17.4% of shortest paths have 1 independent shortest path, 6.0% have 2 independent shortest paths, etc. The final statistics show that 93.7% of shortest paths have at least 1 independent shortest path. This fully demonstrates that the current network has very high path redundancy. However, due to limitations of current Internet routing protocols, including inter-domain BGP and intra-domain OSPF, these large numbers of hidden paths usually cannot be discovered and utilized by routing protocols. For example, many ISPs do not accept BGP announcements from contiguous address blocks smaller than /8192. Thus, even if these networks have multiple connections with other networks, they become hidden paths that cannot be used because they cannot be found in BGP routing tables. Additionally, with network development, many private links have emerged between networks. These private links could originally be used to optimize transmission but cannot be found by current network protocols. Overlays can discover and utilize these redundant paths through active probing and interaction between nodes.

3.3.2 Improvement Capability Analysis of Inter-Domain Routing Optimization Overlay Technology The degree of improvement that overlay technology provides over traditional Internet routing in terms of network performance directly relates to its research value and application prospects. Therefore, debates about the practicality of overlays have never ceased. As mentioned earlier, research on overlays in the past decade can be roughly divided into relay-type overlays and virtual network-type overlays. The latter is often related to optimizing specific business indicators and is difficult to quantitatively evaluate and analyze. Therefore, most simulation and testing schemes have focused on relay-type overlay routing.

Test Network Scale	Packet Loss Rate Improvement	Throughput Improvement	System Overhead
a. US RON vs BGP	13 nodes, including 2 European nodes	Over 11% of sampled groups can be improved by more than 40ms	When network is good: 5% packet loss rate; When 10%: over 15% improvement; over 10% improvement
b. US Overlay vs Multi-homing vs BGP routing	68 nodes distributed across 17 US regions	Based on PlantLab s3' s 588 nodes vs Multi-homing: 5%-15% vs BGP: 33% average delay within 100ms	Proportion of numbers improved by over 30% (50%→80%)
c. Japan Overlay routing vs BGP routing	18 nodes including Tokyo, Osaka, and 4 regions	Proportion of numbers within 100ms improved by over 15% (80%→95%)	$O(N^2)$
d. Japan Ping-based overlay routing vs BGP	43 nodes including 22 foreign nodes	Average 17% improvement	$O(N^2)$
e. US PlantLab Indirect routing vs BGP routing	100Mb connections increased by 10% (80%→90%)	$O(N^2)$	
f. US Bandwidth-aware overlay routing vs BGP routing	Based on PlantLab s3' s 174 nodes	33%~49% improvement	$O(N^2)$
		Average over 20% improvement (over 40ms)	$O(N^2)$

Test Network Scale	Packet Loss Rate Improvement	Throughput Improvement	System Overhead
		About 3% of node pairs can provide low delay and high bandwidth	

Table 1: Statistics on Relay-Type Overlay Optimization Performance

As early as 2001, D. Andersen et al. proposed RON, which later became a typical representative of relay-type overlay routing. Their test results showed that the optimization capability of overlay routing is related to the network environment (as shown in Table 1a). Akella, J.M. Opos, and Sung-Ju Lee conducted comparative tests of overlay and BGP routing performance in larger network environments [40][46][49], with optimization effects more significant than those in RON tests (as shown in Table 1b, e, f). Hiraoka and Uchida [47][48] tested the improvement degree of overlays on network delay and throughput in Japan using PlantLab platforms and different measurement schemes (as shown in Table 1c, d). In all test schemes, overlays adopted relay forwarding technology, providing optimal transmission paths for end-to-end services by traversing all forwarding node pairs, so their system overhead is $O(N^2)$.

The same cross-domain optimization overlay scheme, even on the same test platform, shows very large differences in performance improvement. The reason is that Akella and Opos' s test schemes are incomplete tests, meaning they only selected some poorly performing network nodes as test points. This leads to average performance improvements much higher than other complete test schemes. Even with complete test schemes, due to different network environments, Hiraoka and Uchida' s measurements show large differences in delay optimization results. Overlays' delay improvement capability for real node platform communications is not as good as their performance on PlantLab platforms. S.J. Lee' s complete test results show that although overlays have a large number of redundant paths, the number of nodes that can simultaneously provide low delay and high bandwidth forwarding accounts for only a very small proportion. Significant throughput improvement requires sacrificing delay performance.

Comprehensive analysis of these test results reveals that overlay technology provides more obvious system performance improvements in poorer network environments. This is because when the network environment is good, relay-type overlays require one or multiple forwarding operations at the application layer or network layer of intermediate nodes, and the additional overhead generated reduces optimization effects. Additionally, due to the existence of a large number of redundant paths with varying performance, overlays can specifically optimize for a particular metric needed by an application, thereby achieving greater improvement in related performance. Therefore, under current Internet conditions where robustness and stability need improvement, researching how to deploy overlay technology at the application layer and solve its key problems obviously has great practical value.

4 Efficient Inter-Domain Routing Optimization Overlay Implementation

As mentioned earlier, since network-layer overlays are generally provided by third-party ISPs, such as Detour and OverQoS, although they have the advantages of good forwarding performance and high stability, their deployment flexibility and scalability defects are also obvious. Moreover, network-layer overlays generally only optimize IP-layer inter-domain traffic and cannot segment terminal application requirements. Considering overlay scalability and deployability requirements, application-layer implemented overlays such as Spines and RON have more development prospects. Terminal applications can use this type of overlay to find transmission paths that truly suit their own needs. For example, online games can select paths with high throughput and low delay; P2P file transfers can select low-cost paths; and for some mission-critical services, more reliable transmission protocols can be built at the upper layer to improve reliability.

However, application-layer overlays also have their own weaknesses. For example, existing applications basically do not consider underlying network structure information, which will lead to performance degradation due to path correlation during transmission optimization and may also cause the aforementioned vertical network conflicts. Second, when application-layer forwarding node selection and networking are performed, fully meshed connection probing is generally used, with overhead at the $O(N^2)$ level (such as in RON), which seriously affects network scalability. Additionally, multi-path routing is an important means for overlays to achieve transmission optimization [50][51], but there are still problems that need in-depth research and solutions, such as how to effectively reduce the impact of link dynamic changes on network transmission, how to suppress routing switching oscillations, how to perform congestion control, and how to deploy forwarding nodes.

4.1 Scalable Overlay System Architecture

Considering deployment cost and feasibility, most practical overlay networks currently use application-layer relay models. Due to the instability of terminal hosts, this approach can easily cause 动荡 changes in overlay logical topology, seriously affecting forwarding performance. However, simply using fixed forwarding nodes deployed in the network to build the network also faces system cost and flexibility issues.

Therefore, scalable inter-domain routing optimization overlay systems should fully utilize the advantages of both types of nodes and adopt a hierarchical division and hybrid approach to network construction. Specifically, pre-deployed service nodes should mainly be used for overlay network layering, management, and maintenance, and undertake some business forwarding functions, while screened terminal hosts undertake main business routing services (as shown in [Figure 7: see original paper]), forming a three-layer structure with ordinary user nodes.

To overcome the difficulties of various overlays in terms of stability, deployment cost, and scalability, this three-layer architecture requires further in-depth research in forwarding node screening technology, incentive mechanisms, and dynamic network management.

4.2 Underlying Information-Aware Path Selection, Logical Networking, and Probing

The main problems limiting existing application-layer overlay performance and scale still lie in the existence of vertical network conflicts. Most overlays basically do not consider underlying network structure and traffic information. They independently maintain their own logical networks and monitor all possible alternative path qualities through flooding to find effective replacement paths. For example, RON attempts to recover path failures or congestion problems as quickly as possible by sending probe packets at fixed time intervals. This small-interval, frequent probing brings huge $O(N^2)$ overhead to the network. Therefore, RON' s node scale cannot exceed 50 nodes.

Nevertheless, in actual networks, about 40%-50% of path failures are still unrecoverable through this application-layer overlay network. This is because overlay backup paths are correlated with default underlying IP path failures. That is, after the physical links or routers shared by overlay paths and underlying IP paths fail, the alternative paths used by the overlay also fail simultaneously. In fact, this is the most fundamental problem that optimized overlay network routing needs to consider.

J. Han' s research [21] shows that randomly selecting overlay nodes without considering underlying structure information will lead to a large number of overlay path overlaps. Therefore, despite using small intervals to probe path performance, overlay networks' ability to quickly respond to failures and congestion is still limited unless overlay paths can guarantee complete separation of IP-layer elements without correlation. This raises the problem of how to construct overlay paths to minimize their correlation with each other or with default paths in the underlying IP network. Cha proposed an algorithm for building overlay paths through incremental deployment of relay nodes [54], with the goal of reducing the number of links coexisting on default paths and overlay paths. This method assumes that intra-domain network topology is known and analyzes and optimizes intra-domain relay node placement based on this. Similarly, the goal of obtaining topology information to select appropriate relay nodes for building overlay paths is also to reduce the number of overlapping network elements between default paths and overlay paths.

4.2.1 Forwarding Node Selection Based on Path Independence Regarding the scalability issues caused by frequent path probing, Rewaskar studied the trade-off between performance benefits and overhead of overlay networks and believed that using methods such as reducing logical connections and extending information exchange intervals could improve the overlay' s performance/overhead ratio [53]. Hasegawa divided overlays into high-density and low-density categories based on whether underlying paths overlap and proposed a method to reduce probing overhead by using superposition of multi-segment paths in high-density overlay networks [56]. The probing scheme proposed by Cheng [55], PPRR (Path Probing Relay Routing), is the most practical. Compared with the $O(N^2)$ overhead caused by traditional overlays' fully meshed connection probing methods such as RON, PPRR' s network overhead is independent of scale. Its principle is that for each overlay node, maintain a "probe set" composed of its relay nodes (such as A, B, C, D, E, F). In each round of probing, the currently best-performing nodes (such as A, B, and C) are placed in the "top set," and all other overlay relay nodes in the "probe set" but outside the "top set" (D, E, and F) are replaced with overlay nodes outside the "probe set" to prepare for the next round of path probing.

Experiments have proven that this scheme can achieve performance close to RON while reducing probing overhead. However, this scheme still has obvious defects: first, although the number of nodes probed each time decreases, each node still needs to maintain a list containing all network nodes, which affects scalability in large-scale networks; second, due to reduced probing scope, some path failure information cannot be obtained in time; third, the selection of probed overlay nodes does not consider path correlation issues, leading to inefficient path replacement.

To address these problems and achieve efficient probing of backup overlay paths to improve overlay network scalability, we believe that an on-demand probing method combined with Landmark technology can be used. The system can establish n landmark nodes. When a new node joins, it performs network ranging with these n landmark nodes, and simultaneously selects points with smaller network distances in the ranging space as forwarding nodes. Specifically, considering only delay metrics, nodes with smaller delay similarity AvD can be selected by calculating:

$$AvD = \frac{1}{n}\sum_{i=1}^{n}|L_{A_i} - L_{B_i}|$$

where L_{A_i} and L_{B_i} are the network distances from user nodes and forwarding nodes to the i-th landmark node, respectively. Simultaneously, using the n landmark points as target nodes to measure the path overlap of forwarding node B.

Using these two metrics, a Logical Forward Network (LFN) is established for users. Finally, when communication is established, end nodes perform final performance probing on relays in the logical forward network and perform path switching or multi-path routing based on real-time performance.

This "on-demand" path probing method completely avoids unnecessary path probing overhead. Through probing of landmark nodes and screening by path overlap metrics, the static construction of the logical forward network is completed. Since the logical forward network already satisfies many overlay forwarding optimization requirements during construction, user nodes only need to perform final screening according to current communication needs, effectively reducing the scope of path probing. Moreover, when default paths fail or performance degrades, it can greatly improve the performance improvement effect of backup paths.

4.3 Multi-Path Routing Technology in Cross-Domain Overlays

Existing IP routing protocols use single-path data forwarding between sources and destinations. The defect of single-path routing is that data throughput is limited by existing routing strategies. This means that even if alternative high-bandwidth paths exist, data may still be transmitted on low-bandwidth paths due to policy restrictions. Moreover, single-path routing is not suitable for wireless ad-hoc networks, where high routing failure rates can cause data transmission fragility.

A novel solution is to develop multi-path overlay routing technology. Multi-path routing algorithms in overlays can separate and reassemble source and destination data. This algorithm can improve network throughput, reduce network delay, balance load traffic, and improve link robustness. As shown in [Figure 8: see original paper], between source and destination nodes, if ISPs use traditional single-path BGP routing, even selecting the optimal path, the maximum transmission bandwidth can only reach 20Mb. Using multi-path routing technology and pushing data simultaneously through three routes, the maximum transmission bandwidth can reach 20+8+10=38M.

Multi-path routing can be further combined with layered coding mechanisms to provide higher reliability transmission guarantees for real-time services by adding partial redundant information [50]. However, multi-path routing in cross-domain overlays also has problems that need to be solved, the most important being: calculation of the number of paths needed for transmitting specific data; how to select appropriate paths to provide corresponding QoS guarantees and balance cross-domain load traffic given a topology; how to design efficient multi-path routing protocols to ensure TCP stability; and how to design stable TCP congestion control mechanisms to improve network capacity given a topology structure and multi-path routing algorithm.

4.4 Fault Recovery Oscillation Suppression

Another key technical issue that is easily overlooked in designing efficient and scalable overlays is the system overhead caused by the overlay' s own fault recovery mechanism. To provide reliable services at the application layer, each layer of existing networks has its own error detection and recovery mechanisms. When failures or faults occur, multiple protocol layers will detect them and use their own mechanisms for recovery. Generally, lower-layer mechanisms react quickly, while upper-layer mechanisms provide better recovery services to meet the needs of different users and applications. In overlays, fault recovery faster than BGP routing is provided by tracking and probing along the path.

The probing frequency becomes an important system parameter: lower probing frequency means lower system overhead but requires longer time to detect failures. Higher frequency probing, while causing higher overhead, enables overlays to have faster fault response speed.

When overlay fault recovery mechanisms and underlying fault recovery mechanisms have potential conflicts, such as when traffic congestion and transmission network failures generate many unnecessary fault events, this frequent probing-induced path switching will lead to network instability and seriously affect network performance.

Therefore, certain control strategies need to be implemented for path probing periods and path switching methods to limit path switching oscillations. For example, consider introducing a hysteresis dual-threshold exponential decay algorithm or random backoff mechanism, so that overlays can avoid routing oscillations caused by unstable network performance while also taking into account rapid fault recovery performance.

5 Conclusion and Outlook

Due to the distributed management of the Internet leading to inherent defects in guaranteeing overall transmission quality, inter-domain routing optimization overlays have become a hot technology for solving this problem after innovative technologies and traditional improvement schemes have encountered obstacles. This paper reviews the current state of Internet routing technology, provides a refined classification of existing overlay types, and analyzes in detail the advantages and disadvantages of each type. On this basis, we deeply discuss various problems such as network conflicts and traffic oscillations faced by applying inter-domain routing optimization overlay technology. Simultaneously, through analysis of simulation and experimental data from a large number of overlay research literature, we provide a comprehensive quantitative summary of their routing performance improvements. The results show that in the existing Internet network environment, overlays have relatively obvious effects on routing performance improvement. Moreover, even when multiple overlays coexist, they can still achieve better transmission performance than default routing. Therefore, ISPs and other Internet stakeholders should strengthen their attention to and emphasis on the evolution of overlay technology, and provide good interaction and coordination technologies at the two different layers of routing optimization as much as possible to reduce vertical network conflicts and ensure Internet performance. Finally, we summarize the key technical issues in implementing efficient and scalable inter-domain routing optimization overlays from several aspects, including system structure, path probing, oscillation suppression, and multi-path routing, and provide relevant recommendations.

References

[1] S. Blake, D. Black, and et al. An architecture for Differentiated Services. IETF RFC 2475, 1998.

[2] R. Braden, D. Clark, and S. Shenker. Integrated Services in the internet architecture: an overview. IETF RFC 1633, 1994.

[3] V. Paxson. End-to-End Routing Behavior in the Internet. Proc. ACM SIGCOMM, 1996.

[4] C. Labovitz, R. Malan and F. Jahanian. Internet routing instability. IEEE/ACM Trans Networking, vol. 6, no. 5, pp. 515-558. 1998

[5] W. Li. Inter-domain routing: Problems and solutions. Technical Report TR-128. Department of Computer Science, State University of New York. Feb, 2003

[6] S. Savage et al., "The end-to-end effects of Internet path selection" , Proc. ACM SIGCOMM, pp. 289-99, 1999

[7] C. Labovitz, A. Ahuja, A. Bose and F. Jahanian. Delayed Internet routing convergence. Proc. ACM SIGCOMM. 2000

[8] A. Bremler-Barr et al., "Improved BGP Convergence via Ghost Flushing" , Proc. IEEE INFOCOM,

[9] D. Pei et al., "BGP-RCN: improving BGP convergence through root cause notification," Comput. Netw. ISDN Syst., 48(2):175-94, 2004

[10] W. Xu and J. Rexford, "MIRO: multi-path interdomain routing" , Proc. ACM SIGCOMM, pp. 171-82,

[11] Y. Chen, R. Katz and J. Kubiatowicz. Dynamic Replica Placement for Scalable Content Delivery. Proc. of International Workshop on Peer-to-Peer Systems, 2002.

[12] Y. Chu, S. G. Rao, S. Seshan, H. Zhang. Enabling conferencing applications on the internet using an overlay multicast architecture. Proceedings of ACM SIGCOMM, August, 2001

[13] J. Jannotti, D. K. Giford, K. L. Johnson, et al. Overcast: reliable multicasting with an overlay network. Proceedings of 4th Symposium on Operating Systems Design and Implementation (OSDI ' 00), October

[14] G. Fox. Peer-To-Peer Networks. Web Computing, 2001.5-6.

[15] http://www.planet-lab.org

[16] B. Raman, S. Agarwal, Y. Chen et al. The SAHARA Model for Service Composition Across Multiple Providers. Pervasive Computing, August 2002

[17] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient overlay networks. In Proc. ACM Symp. Operating Syst. Principles, 2001, pp. 131-145.

[18] Zhi Li, Prasant Mohapatra. QRON: QoS-Aware Routing in Overlay Networks. IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, VOL. 22, NO. 1, JANUARY 2004.

[19] Subramanian L, Stoica I, Balakrishnan H, et al. OverQoS: An overlay based architecture for enhancing internet QoS. In: Proceedings of USENIX 1st Symposium on Networked System Design and Implementation (NSDI2004). San Francisco: USENIX Press, 2004. 71-84.

[20] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: Informed Internet routing and transport. IEEE Micro, vol. 19, no. 1, pp. 50–59, Jan. 1999.

[21] J. Han, D. Watson, F. Jahanian. Topology aware overlay networks. IEEE INFOCOM 2005

[22] Z. Li and P. Mohapatra. QRON: QoS-aware routing in overlay networks. IEEE Journal on Selected Areas in Communications, vol. 22, no. 1, pp. 29–40, 2004.

[23] Subramanian L, Stoica I, Balakrishnan H, et al. OverQoS: An overlay based architecture for enhancing internet QoS. In: Proceedings of USENIX 1st Symposium on Networked System Design and Implementation (NSDI 2004). San Francisco: USENIX Press, 2004. 71-84

[24] Z. Duan, Z. L. Zhang and T. Hou. Service Overlay networks: SLAs, QoS and Bandwidth Provisioning. Proc. of 10th IEEE International Conference on Network Protocols (ICNP2002), Paris, France, November 2002.

[25] D. Xu and K. Nahrstedt. Finding Service Paths in a Media Service Proxy Network. Proc. of SPIE/ACM Multimedia Computing and Networking Conference, San Jose, CA, January 2002.

[26] D. Xu, K. Nahrstedt, and D. Wichadakul. QoS and Contention Aware Multi-Resource Reservation. Cluster Computing, the Journal of Networks, Software Tools and Applications, 4(2), Kluwer Academic Publishers, 2001.

[27] Ram Keralapura1;2, Nina Taft2, ChenNee Chuah1, Gianluca Iannaccone2, Can ISPs Take the Heat from Overlay Networks? ACM HotNets Workshop. November, 2004. San Diego

[28] Ram Keralapura, Chen-Nee Chuah, Race Conditions in Coexisting Overlay Networks, IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 16, NO. 1, FEBRUARY 2008

[29] Qiu L, Yang YR, Zhang Y, Shenker S. On selfish routing in Internet-like environments. In: Feldmann A, Zitterbart M, Crowcroft J, Wetherall D, eds. Proc. of the ACM SIGCOMM 2003

[30] A. Soule, A. Nucci, E. Leonardi, R. Cruz, and N. Taft. How to Identify and Estimate the Largest Traffic Matrix Elements in a Dynamic Environment. In ACM SIGMETRICS, June 2004.

[31] T. Roughgarden and E. Tardos. How bad is selfish routing? Journal of ACM, 49(2):236-259, 2002.

[32] R. Feldmann, M. Gairing, Thomas Luecking, Burkhard Monien, Manuel Rode, Selfish Routing in Non-cooperative Networks: A Survey, MFCS 2003

[33] R. Gao, C. Dovrolis, and E. W. Zegura, "Avoiding oscillations due to intelligent route control systems," in Proc. IEEE INFOCOM, 2006

[34] Wenjie Jiang, Dah-Ming Chiu, John C.S. Lui, On the interaction of multiple overlay routing, Performance Evaluation 62 (2005) 229-246

[35] Liu Y, Zhang HG, Gong WB, Towsley DF. On the interaction between overlay routing and underlay routing. In: Makki K, Knightly E, eds. Proc. of the IEEE INFOCOM 2005. Piscataway: IEEE, 2005.

[36] Jessie Hui Wang, Dah Ming Chiu, John C.S. Lui, "A game-theoretic analysis of the implications of overlay network traffic on ISP peering"Computer Networks 52 (2008) 2961-2974

[37] http://www.geni.net

[38] http://cordis.europa.eu/fp7/ict/fire/

[39] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang, "Optimizing cost and performance for multihoming," in SIGCOMM, 2004.

[40] A Akella, J Pang, B Maggs, S Seshan, A Comparison of Overlay Routing and Multihoming Route Control, SIGCOMM' 04, Aug. 30-Sept. 3, 2004, Portland, Oregon, USA.

[41] Yong Zhu, Constantine Dovrolis, and Mostafa Ammar, Combining multihoming with overlay routing, IEEE INFOCOM 2007

[42] Hiroki Okada, Tran Nguyen Trung, Kazuhiko Kinoshita, and Koso Murakami, A Cooperative Routing Method for Multiple Overlay Networks, IEEE 2009

[43] Xiaowei Yang, David Clark, Arthur W. Berger, NIRA: a new inter-domain routing architecture, IEEE/ACM Transactions on Networking (TON) Volume 15, Issue 4 (August 2007) Pages: 775 -788

[44] K. P. Gummadi, H. Madhyastha, S. D. Gribble, H. M. Levy, and D.J. Wetherall, "Improving the reliability of internet paths with one-hop source routing," in Proc. OSDI, 2004, pp. 183-198.

[45] Akihiro Nakao, Larry Peterson, Andy Bavier. A Routing Underlay for Overlay Networks. SIGCOMM' 03, August, 2003

[46] J. M. Opos, S. Ramabhadran, A. Terry, J. Pasquale, A. C. Snoeren, and A. Vahdat. A performance analysis of indirect routing. In Proceedings of the IEEE IPDPS 2007.

[47] Yuichiro Hiraoka, Go Hasegawa, Masayuki Murata, Effectiveness of overlay routing based on delay and bandwidth information, 2007 Australasian Telecommunication Networks and Applications Conference, December 2nd -5th 2007, Christchurch, New Zealand

[48] Masato Uchida, Satoshi Kamei, and Ryoichi Kawahara, Performance Evaluation of QoS-Aware Routing in Overlay Network, ICOIN 2006, LNCS 3961, pp. 925–934, 2006.

[49] Sung-Ju Lee, Sujata Banerjee, Puneet Sharma, Praveen Yalagandula, and Sujoy Basu, Bandwidth-Aware Routing in Overlay Networks, IEEE INFOCOM 2008

[50] J. Apostolopoulos. Reliable video communication over lossy packet networks using multiple state encoding and path diversity. Proceeding of The International Society for Optical Engineering, January 2001, vol. 4310, pp. 392–409.

[51] Z. Cen, M. Mutka, D. Zhu and N. Xi. An Overlay Network Transport Service for Teleoperation Systems. Technical Report MSU-CSE-05-01, Department of Computer Science and Engineering, Michigan State University, 2005

[52] B. Chandra, M. Dahlin, L. Gao and A. Nayate. End-to-end WAN Service Availability. Proc. 3rd USITS, San Francisco, CA, 2001. pp. 97–108.

[53] S. Rewaskar and J. Kaur. Testing the Scalability of Overlay Routing Infrastructures. Proc. of the Passive and Active Measurements Workshop, April, 2004.

[54] M. Cha, S. Moon, C. Park and A. Shaikh. Placing relay nodes for intradomain path diversity. IEEE INFOCOM, 2006.

[55] C. M. Cheng, Y. S. Huan, H. T. Kung, et al. Path probing relay routing for achieving high end-to-end performance. Global Telecommunication Conference (GLOBECOM' 04), Dallas: IEEE Press, 2004: 1359-1365.

[56] Go Hasegawa, Masayuki Murata, Scalable and density-aware measurement strategies for overlay networks, 2009 Fourth International Conference on Internet Monitoring and Protection

Author Biographies:

Li Yanjun: Ph.D., Institute of Computing Technology, Chinese Academy of Sciences

Zhang Guoqing: Ph.D., Master' s Supervisor, Institute of Computing Technology, Chinese Academy of Sciences