

Convergence of Intra-Domain Routing with Centralized Control

Jing Fu¹, Peter Sjödin², and Gunnar Karlsson¹

¹ ACCESS Linnaeus Center, School of Electrical Engineering

² School of Information and Communication Technology

KTH, Royal Institute of Technology

SE-100 44, Stockholm, Sweden

{jing, psj, gk}@kth.se

Abstract. The decentralized control scheme for routing in current IP networks has been questioned, and a centralized routing scheme has been proposed as an alternative. In this paper, we compare the convergence of centralized control scheme with decentralized link-state routing protocols. We first identify the components of the convergence time. Thereafter, we study how to achieve fast routing convergence in networks with centralized control. In particular, we analyze how to distribute forwarding information efficiently. Finally, we perform simulation studies on the convergence time for both real and synthetic network topologies, and study the impact of control element location, link weights, and number of failures on the convergence time. The results show that the centralized control scheme can provide faster routing convergence than link-state routing protocols.

1 Introduction

There are architectural designs for current IP networks which could be reconsidered. While the management of IP networks in an autonomous system (AS) is carried out through centralized servers, control functions, such as routing, are performed in a decentralized fashion on individual routers.

A centralized routing scheme simplifies the management and control processes, and reduces the functions required in the routers. As the architectural advantages and challenges of centralized control are well discussed in research papers [1] [2] [3], we focus on studying the intra-domain routing convergence upon link failures. We compare the convergence process in the centralized control scheme with decentralized link-state routing protocols, and identify the components of the convergence time in both cases. Thereafter, we present how to achieve fast routing convergence in networks with centralized control. In particular, we analyze how to distribute FIBs from the control element (CE) to the routers. We provide a thorough analysis of the number of updated FIB entries upon single link failures, and propose an efficient scheme to distribute the forwarding information. Finally, we present simulation results on convergence time for both real and synthetic topologies for varying network sizes. Also, we investigate the impact of control element location, link weights, and number of failures on the convergence time.

The rest of the paper is organized as follows. Section 2 overviews the related work. Section 3 presents the convergence process in IP networks, for both the decentralized link-state routing protocols and the centralized control scheme. Section 4 studies how to perform FIB updates, including how to distribute forwarding information efficiently. Section 5 shows our simulation setup. Section 6 presents and analyzes the results. Finally Section 7 concludes the paper.

2 Related Work

The importance of network management and control has been well recognized by research communities. In particular, there has been research initiatives focusing on centralized control in IP networks recently. The position papers [1] [2] advocate a 4D architecture, which decompose the network operation into *decision*, *dissemination*, *discovery*, and *data* planes. The 4D architecture moves the control functions from individual routers to a centralized server, which enables direct control of IP networks. The evolution of IP networks and distributed router architectures has been reviewed in [3]. In particular, the centralization of control functions has been identified as a logical next step. Eventually, an entire AS may be considered as a distributed router.

Based on the vision of the 4D system architecture, a route control platform (RCP) has been developed to control inter-domain routing centrally [4] [5]. RCP computes BGP routing table for an entire AS in a centralized server, and distributes the BGP routing table to individual routers. However, the intra-domain routing is still performed on individual routers.

The convergence of intra-domain routing protocols has been well studied. The measurement results for the components of the convergence time in OSPF have been presented [6]. Also, a variety of studies have been performed to reduce the convergence time in OSPF and IS-IS [7] [8] [9].

3 Convergence in IP networks

In this section, we first summarize the convergence process in decentralized link-state routing protocols. Thereafter, we present how to achieve fast convergence in networks with centralized control. Finally, we analyze the components of the convergence time.

3.1 Convergence in OSPF

OSPF and IS-IS are the link state intra-domain protocols used in IP networks today. Since the convergence processes of OSPF and IS-IS are identical, we use the OSPF terminology in the rest of the paper for ease of reading.

The convergence process of OSPF after a link failure can be summarized as follows. First, each router exchanges *hello* messages with its neighboring routers to determine its local topology; a link failure can be detected after loss of three consecutive *hello* messages. When a link failure is detected, the router generates a link-state advertisement (LSA) describing the topology change. The LSA is flooded reliably to all neighboring routers. Upon reception of an LSA, it is reflooded to all routers except to the routers

connected to the interface where the LSA arrived. After reflooding, the router recomputes its shortest path tree (SPT) and updates the FIB by merging the SPT and the BGP routing table. The components of the convergence time are presented in the list below.

- Failure detection time
- LSA generation time
- LSA flooding time
- SPT computation time
- FIB update time

3.2 Convergence in networks with centralized control

In the centralized control scheme, a link failure can be detected in the same way as in OSPF. After that, the router reports the failure to the CE by sending a unicast LSA. When the LSA reaches the CE, it recomputes the SPTs for all routers in the network. Thereafter, the CE updates the FIBs by generating and distributing them to the routers. Each router receiving the updated FIB installs it on the line cards. The components of the convergence time are listed below.

- Failure detection time
- LSA generation time
- LSA reporting time
- SPT computation time
- FIB update time

3.3 Components of convergence time

The comparison of the time components in both controls schemes show that the *failure detection time* and *LSA generation time* are similar while there are differences between the *LSA flooding time* and *LSA reporting time*. In addition, the *SPT computation* and *FIB updates* are performed differently in the two control schemes.

Failure detection time In OSPF, the default *hello* timer is 10 seconds, which means it requires 40 seconds to detect a link failure. Recent work suggests that the hello timer can be set to the sub-second range to provide faster failure detection [7] [9].

LSA generation time After a link failure is detected, an LSA is generated by the router. The generation time upon a single failure is fairly short, between 4 ms and 12 ms [9].

LSA flooding time and reporting time In OSPF, the flooding time of an LSA depends on the link latencies and LSA processing time. The latency of a link is approximately proportional to the distance between the routers, and can be between 10 μ s to 66 ms according to the measurement of a real ISP topology, AS 1239, the Sprint network [10]. The LSA processing time is between 2 and 12 ms as measured in [6] [9]. In addition, a pacing delay of 33 ms is introduced to specify the minimum time interval between sending two consecutive LSAs through the same interface.

In the centralized control scheme, the time for an LSA to reach the CE depends on the link latencies along the path and the forwarding delay in each router. The forwarding delay is in the range of tenths of microseconds and is negligible.

SPT computation time In OSPF, each router needs to determine its SPT, which is a *single-source shortest path problem*. In the centralized control scheme, the CE needs to determine the SPTs for all routers, which is an *all-pairs shortest path problem* since it has to find shortest paths between all pairs of nodes. In addition to Dijkstra’s shortest path first (SPF) algorithm [11], incremental SPF (iSPF) that constructs a new SPT based on the previous SPT are available [12] [13]. Using iSPF in a network with 700 nodes, the average computation times are 3 μ s and 2 ms respectively for single-source and all-source problems [13]. Finally, a parameter *SPF delay* specifies the time between a new LSA arrives and the start of the SPF computation, which by defaults is 5 seconds. The parameter is set to reduce the amount for computation in routers and route oscillation.

4 FIB Updates

In OSPF, the FIB updates are performed in the routers and consist of FIB generation and FIB installation. While in the centralized control scheme, in addition to FIB generation and FIB installation, FIBs need to be distributed to the routers.

4.1 FIB generation

The FIB generation process is similar in both OSPF and the centralized control scheme, except that the FIBs are generated in the routers in OSPF, and in the CE in the centralized control scheme.

The FIB generation is to merge the computed SPT in OSPF and the BGP routing table. A BGP route entry contains a number of attributes such as prefix, nexthop, and AS-path. When an AS receives a BGP route from another AS, the BGP nexthop attribute is set to the IP-address of the border router which announced the route. All packets matching a prefix are forwarded to the specified BGP nexthop according to the route entry. However, not all routers are linked directly to the border router, therefore, each router needs to determine the shortest path to the border router by consulting the SPT. The FIB generation process is simply to replace the BGP nexthop with a direct nexthop, which routers can reach directly. The FIB generation time is in the ranges of milliseconds according to our implementation.

4.2 FIB installation

After a FIB is generated, it is installed on the router’s line cards. The FIB installation time is heavily dependent on the technology used in the line cards. For ternary content addressable memory (TCAM) based line cards, the installation time is linearly dependent on the number of updated entries. Recent results show that a TCAM can perform 1.3 million updates per second [14]. Therefore, updating a FIB entry requires about 0.77 μ s.

Current BGP tables contain approximately 160000 entries [15]. The number of OSPF entries is approximately equal to the number of links in the AS, which is between tens to thousands of entries depending on the network size. The number of FIB entries is the sum of the number of BGP entries and OSPF entries. Since each entry requires 0.77 μ s to update, the total update time is 123 ms for a FIB of 160000 entries.

Finally, the FIB generation and the FIB installation in routers can be performed in parallel. The FIB generation requires CPU usage and the FIB installation time is mostly for accessing the TCAM, and installation of FIB entries on TCAM can be started without waiting for the entire generation process to complete. Since the FIB generation is faster than the FIB installation, therefore, the FIB generation time can be neglected.

4.3 FIB distribution

FIBs are distributed from the CE to the routers in the centralized control scheme. The distribution time is the sum of the transmission delay and the propagation delay.

The transmission delay of the generated FIBs depends on the size of the FIBs and the link capacity. A FIB entry requires 6 bytes: 4 bytes for the prefix, 5 bits for the prefix length, and 11 bits to represent 2048 potential next hops. Therefore, a FIB of 160000 entries is about 940 kB in total. When transmitting the entire FIB on a 1 Gb/s link, the transmission delay is approximately 8 ms. The total time of transmitting all FIBs to the routers is the sum of times for transmitting individual FIBs. However, if there are multiple links connected to the CE, multiple FIB transmissions can be performed in parallel, reducing the total transmission delay by a factor of E , where E is the number of links connected to the CE. In a network of 1000 nodes and 4 links, the total transmission delay is about 2 seconds, which can be a dominant factor in the convergence time.

4.4 Number of updated FIB entries

The discussion on FIB distribution in the previous section assumes that the entire FIB is distributed. However, the number of updated FIB entries upon a link failure could vary significantly. In addition, upon a link failure, the number of updated FIB entries varies for each router in the network.

To have a realistic understanding of how many FIB entries are updated upon a single link failure, we have performed a study on the Swedish University Network (SUNET) which consists of 54 routers [16]. We have generated 20 random link failures, by comparing the FIBs before and after each link failure, we obtain the number of updated FIB entries in each router. The sum of these numbers from all routers is the total number of updated FIB entries in the entire network, which is the total number of FIB entries the CE needs to distribute.

Fig. 1 shows the percentage of updated FIB entries in the entire network upon each failure, sorted in increasing order. The figure shows that for SUNET routers, about half of the link failures have no large impact on the FIBs. In addition, at most 11%, and in average 2.2% of the FIB entries are updated. Therefore, in the real case, significant less FIB entries are updated, and transmission delay should be significantly shorter than in our worst-case analysis that assumes distribution of full FIBs for all routers.

4.5 An alternative approach to FIB distribution

Although we have shown that at most 11% of the FIB entries need to be distributed upon a single link failure, the transmission delay could still be high for large networks. 11% of the 2 seconds required for distribution of full FIBs in networks with 1000 nodes are still 220 ms, therefore, alternative approaches to distribute FIBs are desired.




Fig. 1. Percentage of FIB entries to be updated in the network upon single link failure, 20 random link failure locations.

The separation of intra-domain and inter-domain routing has been proposed in [4]. The separation simplifies the routing process, and is sound from a network architecture point of view. An alternative approach to distribute FIBs is then to maintain two tables, both being distributed from the CE. The first table is a BGP routing table, mapping BGP prefixes to BGP nexthops. The size of this table is equal to the number of BGP prefixes. However, the BGP prefixes do not change upon a local topology change in the network. The second table is an internal table that maps BGP nexthops to direct nexthops, and it also contains intra-domain routes. The number of BGP nexthops is small, about 130 for the SUNET. Combining it with the intra-domain routes, the number of entries is about 500 in total. Although this number grows for larger networks, not all entries need to be updated upon a single link failure. Therefore, we assume that distributing 500 entries is adequate for single link failures.

In this approach, the internal tables are distributed from the CE to the routers upon local topology change. Using the BGP and the internal table, the FIB is constructed in each router. This construction process is similar to the FIB generation process, and is faster than FIB installation on TCAM. Therefore, the construction time can be neglected since construction and installation can be performed in parallel.

5 Simulation Setup

To evaluate the convergence time in various setups, we have developed a Java simulation tool based on the SSFNet simulation framework [17]. The network topologies and the parameters used in the simulations are described and specified in the following subsections.

5.1 Network topologies

We have created several synthetic router-level topologies using the Waxman [18] and the BA [19] models. The number of nodes is varied between 10 and 1000 to demonstrate the scalability. In both models, nodes are placed in a square, with the distance

of each side corresponds to a latency of 30 ms. The link weights have been set either proportionally to link latencies or given unit weights to all links.

In addition to the synthetically generated topologies, we have used two real ISP topologies, which are measured in Rocketfuel based on the *traceroute* technique [20]. The first topology we used is from *AS 1775*, the Ebone network in Europe, which consists of 88 backbone routers and 161 links. The second topology is from *AS 1239*, the Sprint network in the USA which consists of 323 routers and 975 links. The link latencies of the real topologies are set based on the end-to-end measurements performed in [10]. The links have either been given unit weights or given the real weights inferred in [10].

5.2 Providing communication paths in the network

One challenge in the centralized control scheme is to provide robust communication paths between the CE and the routers, in particular upon link and router failures, since the routers rely on the network to report the topology changes, and the CE relies on the network to control the routers. There are several classes of solutions to address this problem: source routing, spanning-tree protocols and flooding. A practical solution used in our simulations is as follows: Upon a link failure, at least one of the routers does not rely on the failed link to reach the CE. Upon receipt of an LSA, the CE determines which routers can be reached; for unreachable routers, source routing is used.

Table 1. Simulation parameter settings, default values in **bold** font.

Name	Value
Network topology	Waxman, BA, AS 1775, AS 1239
Network size	10, 88, 100, 323, 1000
Link bandwidth	1 Gb/s
Link latencies	1-66 ms
Link weights	latency-based, real , uniform
LSA generation time	8 ms
LSA processing time	2, 12 ms
All-pairs SPF computation time	0, 3 ms
LSA pacing delay	33 ms
SPF delay	5 seconds
Routing entries to distribute	500
FIB entries to update	160k
FIB entry update time	0.8 μ s
Number of failures	1 , 5

5.3 Parameter settings

Table 1 summarizes the parameters used in the simulations, which include network topology and protocol-related parameters. The topology-related parameters are described in the previous subsection, and the protocol-related parameters are derived from the time components discussed in Sections 3 and 4. In particular, we have used the approach of maintaining two tables on routers as explained in Section 4.5, and 500 entries

are distributed upon a link failure. Also, we have excluded the failure detection time from the convergence time since it is the same in both control schemes.

Parameters with multiple values have been set to explore their impact on the convergence time. The default values for these parameters are shown using **bold** fonts. For link weight settings, latency-based weight setting is the default for synthetic topologies and real link weight setting is the default for real ISP topologies. For the SPF computation time, we assume it is 3 ms for all-pairs SPF computation for networks of 1000 nodes, while for other setups, we neglect the computation time since it is below 1 ms.

6 Results

6.1 Convergence time with respect to CE location

The convergence time after a link failure in the centralized control scheme depends on the location of the CE. The CE location can be manually chosen by the network operator to reduce the convergence time. To find a reasonably good CE location, we have studied ten different CE locations in each network topology, and performed 50 simulations for each CE location to study the convergence time. The candlesticks in Fig. 2 show the min, max 5th and 95th percentile, and median convergence time for a given CE location. As can be seen in the figure, the convergence time varies depending on the CE location, and for most of the topologies, there are clearly better and worse CE locations which can be identified. The comparison of convergence time for the network topologies shows that larger networks lead to higher convergence time in general. Also, the convergence time of the Waxman topology of 100 nodes and the BA topology of 100 nodes are similar.

In the figure, all convergence times are above 5 seconds, due to the SPF delay of 5 seconds imposed by the CE according to the OSPF specification. When lower SPF delay is set to achieve faster routing convergence, our results are still valid, with only the offset at y-axis in the figure to be reduced accordingly. Finally, a smaller SPF delay results in a larger difference between the best and the worst convergence time in percentage, which in turn increases the importance of a good CE location.

6.2 Convergence time comparison with OSPF

Using the best CE locations identified in the previous experiments, we compare the convergence time in the centralized control scheme with OSPF. In OSPF, the convergence time does not depend on link weights since it is based on flooding. However, in the centralized control scheme, appropriately set link weights reduce the time to report the failure and reduce the propagation delay when distributing FIBs.

Fig. 3 shows the convergence time for OSPF and the centralized control scheme. Each curve in the figure represents the convergence time for 50 simulations upon a single link failure at random locations, sorted in increasing order.

The results for synthetical topologies in Fig. 3(a, b, c) show that the convergence time in OSPF is shorter than the time in the centralized control scheme if uniform link weights are used, however, when latency-based link weights are used, the centralized

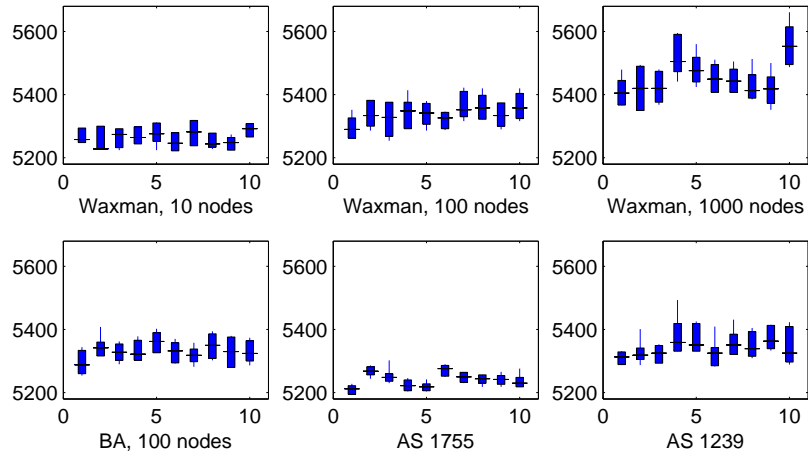


Fig. 2. Impact of CE location on convergence time for various network topologies. X-axis shows different CE locations, and y-axis shows the convergence time in ms. Each candlestick represents min, max, 5th and 95th percentile, and median convergence time for a given CE location. (5000 ms SPF delay is added as it is imposed by the standard)

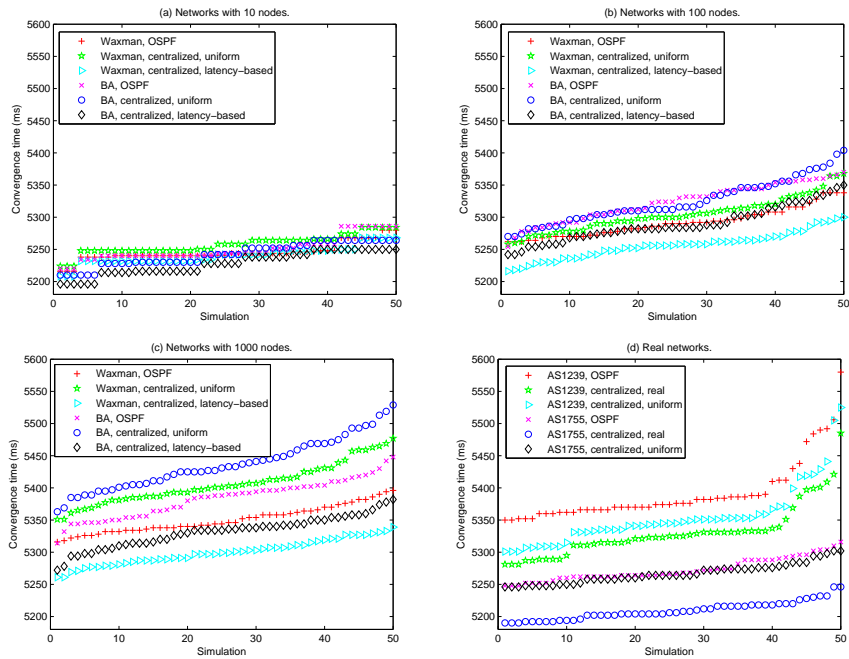


Fig. 3. A comparison of convergence time in OSPF and the centralized control scheme, for both latency-based, real, and uniform link weights. Each curve represents 50 independent simulations with convergence time sorted in increasing order.

scheme has a shorter convergence time. This particularly highlights that fast communication paths between the CE and the routers are important for fast convergence. The results for real topologies in Fig. 3(d) show that the convergence time for the centralized control scheme is shorter than for OSPF, specially for real link weights. The difference in convergence time between using uniform and real link weights is about 40 ms.

The reasons for higher convergence time in OSPF than the centralized control scheme are as follows: First, in the centralized control scheme, only a small table of 500 entries are distributed upon a single link failure, which reduces the transmission delay significantly. Second, the fast incremental all-pairs shortest path algorithm reduces the computation time in the centralized control scheme. Third, the LSA processing delay of 12 ms results in higher convergence time in OSPF.

6.3 Convergence time with varying LSA processing time

The LSA processing time varies for different routers and software applications, and future architectures and technologies may result in shorter processing time. Therefore, we study its impact on the convergence time.

Fig. 4 shows the convergence time with varying LSA processing time. As can be seen in the figure, when the LSA processing time is at 12 ms, the centralized control scheme performs better than OSPF. However, when the LSA processing time is at 2 ms, OSPF has lower convergence time. During the flooding process, each router needs to process the LSA. While in the centralized control scheme, the LSA is only processed once in the CE. Therefore, a shorter LSA processing time has a larger positive impact on the convergence time in OSPF.




Fig. 4. Impact of LSA processing time on the convergence time, Waxman topology of 100 nodes.




Fig. 5. Impact of multiple link failures on the convergence time, Waxman topology of 100 nodes.

6.4 Convergence time upon multiple link failures

Multiple link failures may occur simultaneously in the network, therefore, we present the convergence time upon multiple link failures for both control schemes in Fig. 5.

As can be seen in the figure, for single link failures, the difference on convergence time between OSPF and the centralized control scheme is about 40 ms. However, when there are five *simultaneous* link failures, the difference on convergence time becomes larger, to about 100 ms. This indicates that the centralized control scheme is more resilient against multiple link failures. In OSPF, each router detecting the failure floods out an LSA, and a pacing delay of 33 ms is applied in routers when reflooding two consecutive LSAs, which increases the convergence time. While in the centralized control case, the extra convergence time upon multiple link failures is equal to the difference between the latest LSA arrival time at the CE and the LSA arrival time in the single link failure case.

7 Conclusions

In this work, we have studied the convergence process of intra-domain routing with centralized control, and compared it with the convergence process in OSPF. In particular, we have identified the components of the convergence time.

Our analysis shows that in the centralized control scheme, the distribution of full FIBs of 160000 entries to the routers does not scale well, therefore, we have investigated the total number of updated FIB entries upon single link failures. Our results show that at most 11%, and in average 2.2% of the FIB entries need to be updated. Therefore, faster FIB distribution can be achieved by distributing only updated entries. In addition, we came up with the approach of maintaining two separate lookup tables for FIB updates. Upon a link failure, only a small table of 500 entries needs to be distributed to the routers, which significantly reduces the FIB distribution time.

We have also studied the impact of CE location, link weights, LSA processing time, and number of link failures on the convergence time. The results show that there are clear better and worse CE locations in the centralized control scheme, therefore, a good CE location should be chosen to minimize the convergence time. The results also show that latency based link weights reduce the convergence time for the centralized control scheme, since both the reporting of the failure and distribution of the FIBs rely on fast routing paths. The study on LSA processing time shows that a short time is more in favor of OSPF, due to it is processed in all routers during the flooding process. For multiple link failures, the study also shows that the convergence time in OSPF increases significantly more than the time in the centralized control scheme, due to the pacing delay.

Finally, the comparison of centralized control scheme and OSPF shows that a well-designed centralized control scheme can support networks of 1000 nodes, and may provide faster convergence time than OSPF. As the centralized control scheme has architectural advantages, and its purpose is not primarily to reduce the convergence time in routing, our study that shows a comparable convergence time makes the centralized control scheme possible in future IP networks.

References

1. J. Rexford, A. Greenberg, G. Hjalmytsson, D.A. Maltz, A. Myers, G. Xie, J. Zhan, and H. Zhang, "Network-wide decision making: Toward a wafer-thin control plane," in Proc. of *ACM*

- SIGCOMM HotNets Workshop*, November 2004.
2. A. Greenberg, G. Hjalmysson, D.A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "Clean slate 4D approach to network control and management," *ACM SIGCOMM Computer Communication Review*, Volume 35, Issue 5, October 2005.
 3. András Császár, Gábor Enyedi, Markus Hidell, Gábor Rétvári, and Peter Sjödin, "Converging the evolution of router architectures and IP networks", *IEEE Network Magazine*, Vol. 21, No. 4, 2007.
 4. N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe, "The case for separating routing from routers," in Proc. of *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, August 2004.
 5. M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. Merwe, "Design and implementation of a routing control platform," in Proc. of *Networked Systems Design and Implementation*, May 2005.
 6. A. Shaikh and A. Greenberg. "Experience in black-box ospf measurement," in Proc. of *the First ACM SIGCOMM Workshop on Internet Measurement*, November 2001.
 7. C. Alaettinoglu, V. Jacobson, and H. Yu. "Towards millisecond IGP convergence," *The North American Network Operators' Group (NANOG) 20*, Washington D.C., October, 2000.
 8. A. Basu and J. Riecke. "Stability issues in OSPF routing," in Proc of *ACM SIGCOMM 01*, New York, USA, August, 2001.
 9. P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving sub-second IGP convergence in large IP networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, issue 3, pp.35-44, July, 2005.
 10. R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, "Inferring link weights using end-to-end measurements," in Proc. of *ACM SIGCOMM Internet Measurement Workshop*, 2002.
 11. T. H. Cormen, C.E. Leiserson, R.L. Rivest and C. Stein "Introduction to algorithms, 2nd edition," MIT Press, Cambridge Massachusetts, 2001.
 12. P.Narvaez, K.-Y. Siu, and H.-Y. Tzeng, "New dynamic SPT algorithm based on a ball-and-string model," *IEEE/ACM Trans. Networking*, vol. 9, pp. 706-718, Dec. 2001.
 13. C. Demetrescu and G.F. Italiano, "A new approach to dynamic all pairs shortest paths," in Proc. of *the 35th Annual ACM Symposium on Theory of Computing, (STOC'03)*, 2003.
 14. G. Wang and N. F. Tzeng, "TCAM-Based forwarding engine with minimum independent prefix set (MIPS) for fast updating," in Proc. of *IEEE ICC '06*, June 2006.
 15. Route Views Project, University of Oregon. Available: <http://www.routeviews.org/>.
 16. Swedish University Network (SUNET), Available: <http://www.sunet.se> and <http://stats.sunet.se>.
 17. Renesys. SSFNet, Scalable simulation framework for network models. Available: <http://www.ssfnet.org/>.
 18. B. M. Waxman, "Routing of multipoint connections," *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 9, pp. 1617-1622, Dec. 1988.
 19. A. Barabasi and R. Albert. "Emergence of scaling in random networks," *Science*, Vol 286, 15 October, 1999.
 20. N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in Proc. of *ACM SIGCOMM '02*, August 2002.