

Routing Management in the PSTN and the Internet: A Historical Perspective

Deep Medhi¹

¹University of Missouri–Kansas City, Kansas City, Missouri, USA

Abstract. Two highly visible public communication networks are the public-switched telephone network (PSTN) and the Internet. While they typically provide different services and the basic technologies underneath are different, both networks rely heavily on routing for communication between any two points. In this paper, we present a brief overview of routing mechanisms used in the PSTN and the Internet from a historical perspective. In particular, we discuss the role of management for the different routing mechanisms, where and how one is similar or different from the other, as well as where the management aspect is heading in an inter-networked environment of the PSTN and the Internet where voice over IP (VoIP) services are offered.

1 Introduction

Two highly visible public communication networks are the public-switched telephone (PSTN) and the Internet. The PSTN provides voice services while the Internet provides data-oriented services such as the web and email (although voice communication over the Internet is now possible). In both these networks, information units need to be routed between two points. In the PSTN, the unit of information is a call, while in the Internet the unit is a datagram (packet). Thus, in each network, routing plays a significant role so that information units can be delivered efficiently. To accomplish efficient routing, various supporting management functions also play critical roles.

Routing in both the PSTN and the Internet has evolved over the years; for details on routing, refer to books such as [1, 9, 8, 16, 19, 24]. The focus of this paper, however, is not routing. Rather, it is *management* that is related to routing; for brevity, this will be referred to as *routing management* throughout this article. Simply, routing management means types of management functions needed for the efficient operation of routing, which may require determination of optimal paths. Naturally, along with the evolution of routing, routing management has evolved as well. In recent years, there have been significant activities with regard to voice services over the Internet, commonly referred to as the Voice over IP (VoIP) service. The VoIP service has necessitated interworking of the PSTN and the Internet in terms of routing of a call, and thus, on the management aspect. The scope of the paper is to present the historic context of routing management both for the PSTN and the Internet along with evolution of routing, concluding with

interworking management of the PSTN and the Internet for VoIP routing. The goal is for the reader to understand the evolution path for each network separately, and to be able to draw conclusions on the similarities and differences, and convergence issues.

The rest of the paper is organized as follows: in Section 2, we present routing in the PSTN along with the evolving role of management, including routing under number portability; Section 3 covers routing and management in the Internet; finally, we discuss the management aspect with regard to PSTN-Internet interworking for VoIP services in Section 4.

2 Public Switched Telephone Network

2.1 Hierarchical Routing

Telephone networks have been around for over a century. However, the need for any form of routing did not arise until the 1930s. Until then, essentially point-to-point direct links (trunkgroups) were set up to connect calls between different places and no routing was involved. The need for routing arose for two primary reasons: 1) point-to-point links lead to the N^2 problem, i.e., if there are N nodes in a network, we need $N(N - 1)/2$ directly-connected links, and as more and more cities (with multiple switches) offered telephone services, this problem grew significantly, and 2) it was recognized that some trunkgroups were less utilized compared to others. Thus, if there were any ways to take advantage of this by routing calls through less utilized trunkgroups, capacity expansion could be avoided. Capacity expansion used to be very costly and still is in many cases. There is another impetus to arrive at some form of routing: as the switching technology started to move from old mechanical switches to electro-mechanical switches, the possibility of switching being capitalized to perform some form of routing became more than just a thought.

It is important to note that routing in the telephone network was first performed in an age when neither information storage nor information exchange between switches was possible. At that time, all that could be done is in-band signalling to setup a call; no other information exchange about routes or link status could be communicated. In the absence of any information exchange mechanism between switches, a newly arriving call merely hunts an outgoing available circuit at each switching point; in this environment, an important requirement of routing was to be addressed: *how to avoid looping*? That is, through trunk hunting from one trunkgroup to another, how does an arriving call not revisit switches in a circuitous manner? The technology available at that time only allowed the call set up to be accomplished through a mechanism known as *progressive call control (PCC)*, which forwarded in-band set up signaling messages from one switch to the next. A call control cannot revert back to a switch from where it started in order to attempt a different route. For this environment, the Bell System came up with a mechanism for routing that avoided looping: introduce *hierarchy* among network nodes and still use progressive call control. We next briefly discuss the hierarchical routing architecture.

Hierarchical Routing Architecture

There are five levels defined in the hierarchy that were introduced by the Bell System. At the bottom are the end offices or central office switches, also known as class-5 switches. As we move up, we go from toll switching centers (class-4), to primary switching centers (class-3), to secondary switching centers (class-2), to regional switching centers (class-1). They are connected by *final* trunkgroups or *high usage* trunkgroups in cases where there is enough traffic volume to justify such trunkgroups. Five levels of switching hierarchy are shown in Figure 1.

From a geographic perspective, there is another way to view the network, which is a planar view. This is shown in Figure 2. We can see that part of the network that is under a regional switching center is essentially a tree-based network except for any high usage trunkgroup (marked by a dashed line), which connects a switch under one regional switch to another switch in the same or a different regional switch. The network at the regional switching center level (or the highest level if all five levels are not used) is fully connected.

With the hierarchical structure and in the presence of high usage trunkgroups, multiple alternate paths are available between end offices. Calls here attempt these paths in a fixed, pre-defined order. Recall that routing was accomplished without requiring any information exchange between switches, or in an information-less environment. Furthermore, looping was avoided by carefully defining rules for switching hierarchy and forwarding of calls.

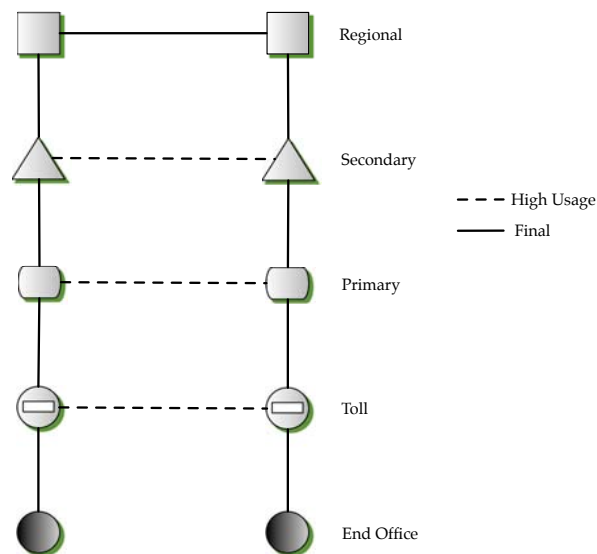


Fig. 1. Switching hierarchy in hierarchical call routing.




Fig. 2. Geographical perspective of hierarchical routing.

Another important issue is how to address routing of international calls from one country to another. In the hierarchical routing structure, another switching level is defined above the regional switching center to connect trunkgroups from one country to another country.

Therefore, hierarchical routing can be briefly summarized in the following way: switches in the network are placed at different levels, and a call can move up using a trunkgroup from a lower level switch to a higher level switch.

Management

As we can see from the structure of the switching hierarchy, the routing path order is pre-determined and the paths are attempted in a certain order. Thus, there is no direct need or use for a management function that determines optimal route selection. On the other hand, for the network to perform well, there are two important aspects to consider along with routing: 1) capacity expansion, and 2) network control. Capacity expansion is needed for traffic growth so that the pre-ordered routes can provide an acceptable grade-of-service; capacity expansion falls primarily under the network planning phase. However, the need for a network management function arises from the perspective of *control* functions so that the network does not perform poorly, especially in an overloaded condition.

There are several possible controls that can be invoked when congestion is detected. In general, controls can be classified as: 1) trunk reservation, 2) dynamic overload control, 3) code control and hard-to-reach (HTR) capability, 4) reroute control, 5) directionalization of a link (trunkgroup), and 6) selective dynamic overload control [18, 40,

34]. Controls are of two types: *restrictive* and *expansive*. Restrictive controls limit traffic from reaching congested location(s); expansive controls allow new paths to explore to avoid congestion. Of the above classifications, only reroute control falls under the expansive category while the rest are restrictive controls.

Trunk Reservation: Trunk reservation [41], also known as state protection [16], refers to logical reservation of part of a capacity on a link (trunkgroup) for its own direct traffic. In effect, trunk reservation depends on a threshold value. If a trunkgroup is not filled with calls before this threshold is reached, a call between other origin-destination pairs can still use a trunk from this trunkgroup; once this threshold is reached, only direct calls, i.e., the calls that originate/terminate at either end of this trunkgroup can make use of this remaining capacity. An advantage of the trunk reservation control is that it stabilizes network behavior, especially in a dynamic call routing environment.

Dynamic Overload Control: Dynamic overload control (DOC) is used for *sensing* switch congestion. When a switch is congested, it makes sense to reduce traffic being directed to this switch by other switches.

Code Control and Hard-to-Reach (HTR) Capability: In this case, instead of looking at the view of congestion from the perspective of a switch, the view is taken at the *destination code level*. Destination codes can be at the central-office level and/or at the actual destination number level. For example, in the North American numbering plan, this means that the code control is activated either at the NPA-NXX or at the NPA-NXX-XXXX level in which NPA represents the area code, NXX is the central office code, and XXXX is the number assigned to the subscriber. Any code control information for such a destination code needs to be pushed to the entry points into the network, i.e., where calls originate. This way, a newly arriving call to such a destination code may not be allowed to enter the network since the hope of establishing this call is very small. This occurs when a switch in a network receives a hard-to-reach code from another switch and it employs code control. There are two forms of code control: *call gapping* and *code blocking*. Call gapping means that calls made to a particular destination are *gapped* every Δt sec; that is, only one call is allowed to pass through every Δt sec. Code blocking allows calls to go through on a percentage basis. While conceptually they sound similar, the difference is that with code blocking, two or more back-to-back calls may still be allowed to go through since probability is used; in a congested situation, this may not be desirable. However, with call gapping, allowed calls are paced evenly.

Reroute Control: This is an expansive control that is used for overflowing traffic to a new route/trunkgroup. It may be noted that traffic in the network is not uniformly loaded; thus, it is not unusual to find underloaded trunkgroups in the network. Through the reroute control, traffic is routed through such trunkgroups as long as the hierarchical routing rule allows it. Note that reroute controls can be activated in a code-specific basis as well.

Directionalization of a Link: Trunkgroups in a circuit-switched telephone network are bi-directional. With a focused overload situation, traffic from one end can be excessive compared to the other, which can increase overall blocking. In this case, it is desirable to let some calls out from the low-traffic end. Thus, *directional* trunk reservation is invoked from the high traffic side to the low traffic side, but there is no trunk reservation from the low traffic side to the high traffic side.

Selective Dynamic Overload Control: Dynamic overload control (DOC), which is described earlier in this section, acts at a generic level; it cannot distinguish between codes with high-completion probability and codes with low completion probability. SDOC is a merger of DOC and code blocking and therefore, traffic cancellation can be activated for HTR codes.

To summarize, above control mechanisms work in concert with routing to provide optimal network performance management under a variety of conditions including under overloaded situations. Note that these schemes often require information to be exchanged between switching nodes. Since they are data-oriented information, they are communicated through an out-of-band mechanism; this is commonly done using the SS7 signaling network. In many instances, such exchanges may involve multiple service providers. For example, suppose that a HTR code is activated for a destination code point in a provider's network; this information may need to be communicated to other providers so that these providers can activate call gapping in order to reduce calls destined for HTR codes from entering at the very edge of the network. This means that at the SS7 signaling network level, coordination is also necessary to pass information from one provider to another.

2.2 Dynamic Call Routing and Its Management

Starting in the mid 1970's, the concept of dynamic call routing was investigated in order to determine ways to address inflexibilities of hierarchical routing. For example, due to the hierarchical nature, a series of certain trunkgroups may not be fully utilized between two endpoints, even though they have idle trunks. Secondly, busy hour traffic may not coincide due to different time zones. For example, in the continental United States, non-coincidence of busy hour traffic cannot be exploited by fixed hierarchical routing.

Several different dynamic routing concepts have been developed and deployed such as dynamic nonhierarchical routing (DNHR) [2], real-time network routing (RTNR) [4], dynamically controlled routing (DCR) [38], and dynamic alternate routing (DAR) [15]. Briefly, the main idea behind all dynamic routing schemes was to have flexible alternate paths that can exploit time-based load, or load change dynamically/adaptively, without being limited by hierarchy rules; for details, see [24]. A historic view of deployment of dynamic call routing can be found in [3]. These dynamic routing schemes were deployed above the class-5 access switches; that is, a class-5 switch serves as the entry

point and still can be thought of as the lower end of the hierarchy; the switches at class-4 and above levels were combined into a flat dynamic routing architecture.

Briefly, we next discuss each routing scheme and the role played by routing management. DNHR is the first implemented dynamic routing scheme. It was developed by AT&T and was deployed in AT&T's long-distance telephone network in 1984; it was later retired in 1991 when RTNR was deployed. We discuss it here primarily for its historical context and especially with regard to the management aspect that arose with its operation.

DNHR is a time-dependent routing scheme. This means that the set of routing paths available (and their order) is different at a different time of the day. In the case of DNHR, the 24-hour time period spanning a 7-day week was divided into 15 load set periods (LSPs): 10 for weekdays and 5 for weekends. The different number of LSPs was determined based on understanding traffic patterns. For example, the same routing pattern can be used from midnight to 8 o'clock in the morning due to low traffic volume.

Because of the above requirement, a management issue that arose was the need for a route computation function off-line. This function was to be invoked every week or every other week; the computed routing tables were then uploaded to the switches. Typically, dedicated circuits were used for loading such computed routes. This also meant that the off-line process was required to receive traffic forecasting estimates based on network measurements and relevant business drivers, which were used for route computation. Thus, a critical management system was introduced that can be best described as a routing administration system. In addition, in order to accommodate sudden traffic change in the network, traffic measurements every five minutes were collected to compute on-demand routes which were appended through a network management and operations system (NeMOS).

Note that going from hierarchical routing to DNHR improved network performance considerably due to routing flexibility; in addition, capital cost for network capacity was reduced since significantly less capacity was needed to maintain the same grade-of-service than in a hierarchical environment [1]. However, this came at the expense of requiring a network traffic measurement and management system so that dynamic routing could be effectively exploited; thus, the role of management increased with the introduction of dynamic call routing.

RTNR was introduced in the AT&T network in 1991 and is still in use today. In this scheme, the routing table was determined at the switch on a per call basis. This means that it does not require an off-line computation mechanism unlike DNHR. Thus, the off-line management system of DNHR has been essentially avoided with RTNR; not only that, the adaptive feature of RTNR resulted in better network performance than DNHR. Nevertheless, a real-time management system to monitor trunk status and reporting status in response to a query from a switch would be still needed.

Next, we briefly discuss management issues with regard to DCR. DCR has been deployed in several networks such as Canada's Stentor long-distance network, and the inter-exchange carrier networks of both MCI and Sprint. It may be noted that DCR re-

quires a centralized processor to compute the routing table on a periodic basis, which can be as frequently as every 10 sec. It requires the trunkgroup utilization status to be sent to the central processor to compute routing tables for all switches. Thus, a management function is used with regard to status updates and centralized routing table computation for smooth functioning of DCR. It is worth noting that the central processor can possibly fail or the communication circuits from switches to the central processor can be down, despite efforts through redundancy; thus, DCR has a mechanism to fall back to two-level hierarchical routing if switches cannot obtain any updated information from the central processor.

DAR has been deployed in British Telecom's national network. With regard to DAR, however, there is little management need with regard to collecting traffic or estimating traffic, since it is a completely distributed, adaptive scheme that is based on a sticky random principle. However, to accommodate switch maintenance, it might be preferable to indicate ahead of a time if a certain switch should not be used for alternate route selection; to do that, a monitoring function and updating function is needed to communicate such information to all switches.

It may be noted, however, that all dynamic routing schemes also employ the controls discussed earlier; the only exception is the reroute control since this was already handled directly by dynamic call routing itself.

2.3 Multi-Provider Environment and Number Portability

Until the divestiture of the Bell System in 1983, the entire hierarchy of the telephone network was provided by the same telephone service provider (TSP) in the United States. In fact, in many nations across the world, the telephone network is still provided by a single provider. With the breakup of the Bell System in the United States, different TSPs play different roles in carrying a call. A call originates in an access TSP, a "local-exchange carrier (LEC)," where the call starts from the end office. If the call is destined for another user in the same LEC, the call is routed within its network. When a call starts from one LEC and is destined to terminate in another LEC, the call typically goes through an inter-exchange carrier (IXC) before entering the destination LEC. From a routing hierarchy point of view, IXC enters at the level of the primary switching centers.

In most cases, LECs use a two-level fixed hierarchical routing architecture. An IXC can deploy either a fixed hierarchical routing or dynamic routing. Unless a call terminates in a different country, there is usually at most one IXC involved between the access LECs. For a call going from one country to another country, the call may go through the second country's inter-exchange provider or equivalent before reaching the destination address in another access carrier.

It may be noted that with the opening of local markets for multiple providers, there are now competitive LECs who provide service along with the incumbent LEC. Furthermore, there are multiple mobile providers in a local market.

Finally, a major change in the PSTN in the last decade has been the development of the number portability concept, dictated by regulatory bodies. The basic idea behind number portability is that it allows the subscriber of a telephone number to keep his/her phone number while changing the provider from one to another.

The implication of these market or regulatory changes is that a local call may start from one provider but would terminate in another provider. Thus, peering agreements are needed between different providers. Secondly, due to local number portability, providers need to exchange more specific information about a telephone number than just providing information at the central office (NPA-NXX) level. Thus, the originating provider, in which a call originates, needs to determine how to hand-off the call to the destination, beyond its own network. This brings forth the notion of three networks: the originating network in which the call originates, the donor network in which the destination number used to be originally allocated to, and the recipient network in which this destination number currently belongs. There are currently four different routing schemes to accommodate number portability (Figure 3):

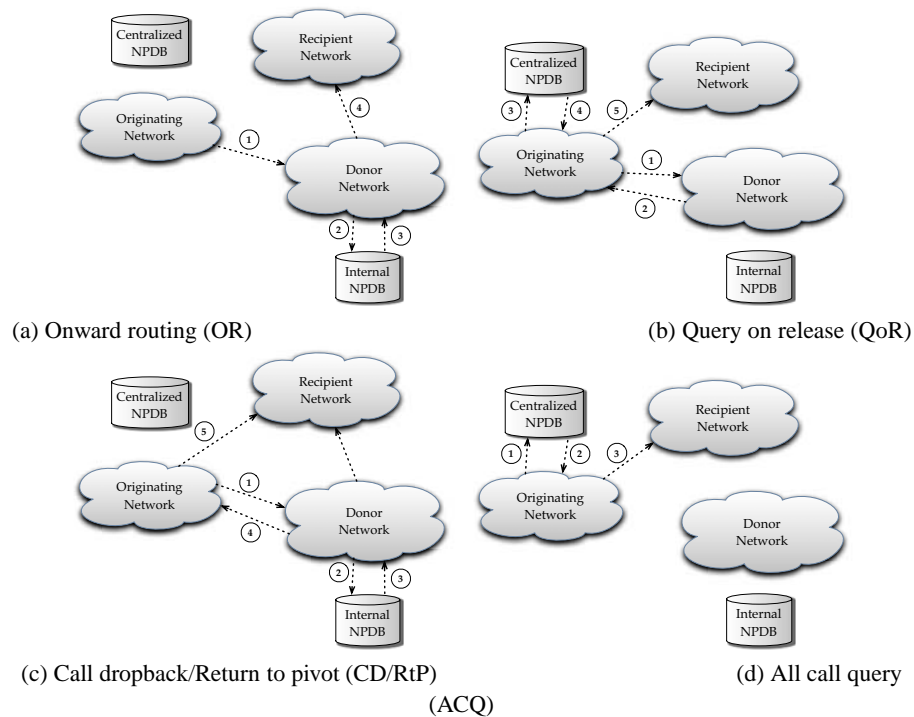


Fig. 3. Call routing schemes for number portability.

Table 1. Comparison of routing schemes for number portability.

Method	Benefits	Drawbacks
OR	1) No centralized database needed, 2) Internal NPDB can be stand-alone and contains only the ported number from the donor network, 3) Good solution for short term, or if a small percentage of subscribers chooses to do number portability	1) Completely relies on the donor network during call set-up, 2) Requires setting up two physical call segments
QoR	1) Centralized number portability database used for call routing decision	1) Involves the donor network during call set-up, 2) Circuits are reserved temporarily to the donor network
CD/RtP	1) Centralized number portability database not needed, 2) Internal NPDB can be stand-alone and contains only ported numbers from the donor network	1) Involves the donor network during call set-up, 2) Circuits are reserved temporarily to the donor network
ACQ	1) Centralized number portability database used for call routing decision, 2) Does not involve the donor network, 3) Efficient in usage of switch ports and circuits, 4) Good long-term solution, especially when most subscribers choose number portability	1) Relatively high portability set-up cost, 2) High ISUP traffic to NPDB from originating switches

- Onward routing (OR): In this case, the call setup message is sent to the donor network as if number portability has not occurred yet for the called number; the donor network is then responsible for forwarding it to the recipient network.
- Query on release (QoR): This also assumes, at first, that the portability has not occurred yet. Thus, the call is routed to the donor network. In this case, however, the donor network informs back to the originating network that the number is ported along with the information about the recipient network. On receiving this information, the originating network establishes the call directly with the recipient network.
- Call dropback/Return to Pivot (CD/RtP): This is a hybrid scheme between onward routing and query on release.
- All Call Query (ACQ): In this scheme, a query is generated to a centralized database regarding the recipient network of a call, regardless of whether the destination number is ported or not.

A comparative summary of these schemes is presented in Table 1. Currently, all four schemes are implemented in practice by different countries [13, 24, 37, 39]. However, within a particular country, typically one of the these schemes is implemented; see the summary in Table 2.

Table 2. Routing methods for number portability currently used in various countries

Country	Scheme
Austria	OR
Belgium	ACQ
Canada	ACQ
Denmark	ACQ
Finland	ACQ
France	OR
Germany	ACQ
Hong Kong	ACQ
Ireland	OR
Italy	OR
Netherlands	ACQ/QoR
Norway	OR, ACQ
Japan	OR
Singapore	OR
Spain	ACQ/QoR
Sweden	ACQ/QoR
Switzerland	QoR
UK	OR
US	ACQ

As we can see from the above discussion and Table 1, a number of management functions are needed for properly routing a call. For instance: 1) database functions are needed that may contain and update portability data periodically, 2) signaling network connectivity to access such databases, for example, through the SS7 network, 3) maintenance of a peering relation database.

In addition, it is possible that there are some “middle-man” networks which are willing to be transit carriers for recipient networks; not only that, but a provider might have peering and economic relation with multiple such transit carriers. Thus, when a call originates, the originating network might make the decision on routing the call to one of its multiply-connected transit networks based on pricing. That is, the role of pricing in economic routing is starting to become a factor, which requires maintenance and exchange of pricing information through an appropriate management system.

3 Internet

3.1 Internet Routing Evolution

We first briefly discuss the evolution of Internet architecture from a historical perspective. Note that we focus on Internet *routing* rather than the Internet as a whole; for an excellent summary on Internet history, refer to [21].

Until the early 1980s, the ARPANET served the role of interconnecting various sites with a rigid two-level hierarchy where the ARPANET nodes were at the top level. In 1983, ARPANET was split, resulting in two networks: ARPANET and MILNET. This was the birth of the two separate networks talking to each other in case one host in one network wants to communicate with another host in the other network, and vice versa. This also resulted in the *need* to have a mechanism by which separate networks could talk to each other. Here “separate networks” means that they are run by different entities.

Exterior Gateway Protocol (EGP), developed in 1982-1984 (refer to RFC 827 [35], RFC 888 [36], RFC 904 [25]) describes how separate networks that are autonomous can talk to each other. Along with EGP, the term *autonomous system* and the notion of a 16-bit *autonomous system number* were introduced in [35]. Briefly, EGP defined a rigid two-level hierarchy with the top level labeled as the core backbone and the bottom level being the level at which the different networks, identified by AS numbers, were connected. NSFNET, deployed first in 1984, relied on EGP. The architecture of and experience with NSFNET and EGP have been documented in [6, 30].

It is important to note that in EGP, non-backbone ASes were not allowed to be directly connected; this is a direct consequence of the strict two-level hierarchy imposed by EGP. Another consequence was that the structure allowed only a single provider at the top level, i.e., the NSFNET. Furthermore, unlike BGP, EGP messages were sent directly over IP without invoking any reliable transport protocol. Thus, if the exchange of information required a large message to be generated, this needed to be handled by fragmentation and reassembly at the application layer of the TCP/IP protocol stack.

In essence, while EGP provided a much needed transitional platform to go from the ARPANET to the NSFNET, it had several restrictions which were not desirable for longer term growth. For example, EGP did not allow ASes to be directly connected. Thus, a network that is located in an AS would need to go through the top level, i.e., the NSFNET, to reach another network in another AS. However, NSFNET faced the situation that some networks that belonged to different ASes had backdoor connectivity; in other words, these networks connected with each other directly without sending traffic via the backbone. Thus, EGP’s strict requirement could not be directly applied nor enforced in the NSFNET. It may be noted that to circumvent the limitation of EGP, a formal confederation approach was suggested in RFC 975 [26]. An important lesson learned from NSFNET in regard to the routing architecture was that no single entity would be managing the global Internet. Each system that is a component of the global Internet will have its own routing paradigm that can be driven by economics and other factors; each such system would have its own interest to connect to other systems directly, instead of using a global core such as the one suggested by EGP. As a corollary, global consensus from the deployment point of view is hard to arrive at; however, a mutual bilateral agreement is possible. Since an AS cannot control what an upstream AS announces, it became necessary to take a policy-driven approach; for example, how the routing mechanism was to handle packets from certain networks based on an import

Table 3. Examples of import and export policies at a BGP speaker.

<i>Import Policy</i>
– Do not accept default 0.0.0.0/0 from AS64521.
– Assign 192.168.1.0/24 coming from AS64822 preference to receiving it from AS64521.
– Accept all other IP prefixes.
<i>Export Policy</i>
– Do not propagate default route 0.0.0.0/0 except to internal peers.
– Do not advertise 192.168.1.0/24 to AS64701.
– Assign 172.22.8.0/24 a multi-exit-discriminator metric of 10 when sent to AS64617.

policy of an AS. This then raises the notion of policy-based routing (see Table 3 for examples of import and export policies). It is to be noted that some rudimentary policy-based routing was done so that certain rule checking could be invoked in the NSFNET as noted in RFC 1092 [30].

EGP and particularly, NSFNET experiences led to the recognition that any future routing architecture must be able to handle policy-based routing (see RFC 1102 [7], RFC 1104 [5], RFC 1124 [20]), and any newly developed exterior gateway protocol must have the ability to handle policy decisions. That is, experience and realization served as the impetus to the development of BGP, which was first introduced in 1989 through RFC 1105 [22]. To summarize, BGP tried to address the following issues: 1) avoiding a strict two-level hierarchy like EGP, 2) allowing multiple levels such that any AS has the option to connect to another AS, 3) using TCP for reliable delivery of BGP data, and 4) making policy-based routing possible.

By 1991, BGP was expanded to BGP, version 3 (see RFC 1267 [23]). At about the same time, it was recognized that the implicit address block classification of an IP address under Class-A, Class-B, and especially Class-C, i.e., classful addressing, would cause a significant growth in the routing table entries at core backbone routers; thus, some mechanisms to avoid/minimize assigning address blocks straight at Class-C were needed. This has led to consider address aggregation through supernetting [14], which subsequently led to the development of *classless inter-domain routing (CIDR)*; this aspect was incorporated in BGP, version 4. BGP, version 4, [31, 32] which is currently used in the Internet.

While BGP, version 4 (BGP4) has resulted in several improvements over BGP, version 3, it is clear that use of CIDR was one of the most significant changes that required communicating netmask information to be advertised along with an IP address block

during a BGP announcement; that is, the addressing structure played a critical role in routing.

Finally, it is worth noting that the notion of dividing a network into a hierarchical structure of intra-domain and inter-domain and allowing each intra-domain to define its own routing can be traced back to the OSI routing model developed in the 1980s; see [29] for further details.

3.2 Routing Management

From a routing protocol perspective, IS-IS and OSPF are the most common routing protocols deployed by moderate to large Internet service providers (ISPs) for routing in their own administrative boundaries while BGP is used for the exchange of IP prefix-level route information among providers. An IP prefix (also known as “route”) refers to an IP address block; an IP prefix is denoted in CIDR notation, which identifies the address block along with the netmask separated by a slash, such as 134.193.0.0/16. The Internet routing architecture is made of autonomous systems that are connected to each other and each autonomous system serves as home to a set of IP prefixes. In most cases, there is an one-to-one relation between an autonomous system and an ISP. However, an ISP may have multiple autonomous system numbers, which might have resulted from a merger of different providers. Similarly, a group of independent providers (with each




Fig. 4. Interconnection of ISPs of different tiers: a representative view.

holding a set of IP prefixes) may be served by a single autonomous system number. For any *IP prefix*, located in a particular autonomous system, a path needs to be found from the rest of the Internet; to be able to accomplish that, BGP is used for exchange of reachability information of IP prefixes. Note that it is possible for an IP prefix to “move” from one autonomous system to another; this occurs when the owner of an IP prefix wants to change its ISP and the new ISP is part of a different autonomous system number. In this case, the new provider will announce the new reachability information about this IP prefix. Note that an AS need not be limited to an ISP; it can be assigned to an organization. For simplicity, we will assume that an ISP (or entity) has one AS number. Then from the perspective of ISPs, they can be categorized as different tiered providers. A simple illustration of different tiered providers is shown in Figure 4. Different tiered providers fall into different types such as a core (“default-free”) provider, a transit provider, or an access provider. The default-free providers are referred to as the tier-1 ISPs, and then depending on the relation, the rest are categorized as tier-2 ISPs, tier-3 ISPs, and so on.

While the basic routing protocols remain the same, the issues faced with regard to management can be quite different. While routing protocols for the Internet have evolved, it is commonly believed in the early days of the Internet that no management was necessary since the Internet could self-manage due to routing protocols being adaptive to changes such as a failure. However, the basic problem with the notion of path selection in Internet routing protocols was that it uses the shortest paths. On the other hand, the basic shortest path routing mechanism does not address how traffic load can be factored for efficient traffic engineering. For instance, if link metrics are simply set to hop counts, then parts of the network could remain congested due to the shortest path routing.

To address this situation, capacity on congested parts can possibly be increased in the next planning cycle. Yet simple capacity expansion might not be the most cost-effective mechanism to maintain a certain level of performance guarantee. Thus, over the years, it was felt that optimal weights for links (“link weights”) can perhaps be determined by taking traffic load into account. This need has led to a series of works starting at the turn of the millennium [12, 27, 28] on link weight determination off-line through an optimization process. In order to implement such an off-line system, it became necessary to introduce a sophisticated management system. Thus, this management system collects traffic data from routers to first determine traffic demand loads [10, 11], and then performs an off-line optimization method to determine the link weights (see Figure 5). The computed link weights are, in turn, loaded to the routers so that they can use the routing protocol mechanism to exchange these link weights as link metric information with other routers. When routers receive this link state information, they can compute the shortest paths, which are then traffic aware; thus, this leads to better network utilization and performance.

Another important aspect is that Internet routing is closely tied to IP addressing. Specifically, a path needs to be found to a destination IP prefix for packet forwarding.

The Internet routing architecture is made of autonomous systems which are connected to each other and each autonomous system serves as home to a set of IP prefixes, with the flexibility that an IP prefix can be moved to a different autonomous system over time. Information about IP prefixes is communicated from AS to AS using BGP. Sometimes, however, an AS announces bogus IP prefix information; the BGP protocol does not have any mechanism to catch such bogus communication. Thus, an important management issue for any provider is to track IP prefixes received and determine through other mechanisms if they are valid and that they indeed belong to the autonomous system number announced through BGP.

Thus, from a routing and route management perspective, there are several aspects to consider: 1) traffic monitoring to determine link weight change, which is then used by the shortest path computation module, 2) IP prefix maintenance, 3) policy issues. We discuss these aspects below.

Traffic Monitoring and Traffic Engineering System The role of such a system is to collect measurement data through traffic monitoring to determine the traffic load matrix. Based on this determination, link weights are then computed and disseminated, which is used by each router to determine the traffic-aware shortest-paths. For large tier-1 providers, periodic link weight determination is an important management function. For this, a traffic monitoring and measurement process is required as shown in Figure 5. Often, however, the link weight determination for a large transit provider is complex since such a provider must also consider *early-exit* routing; when certain destinations have multiple egress points, the issue of choosing the earliest exit point arises, which must be taken into account in link weight determination [33].

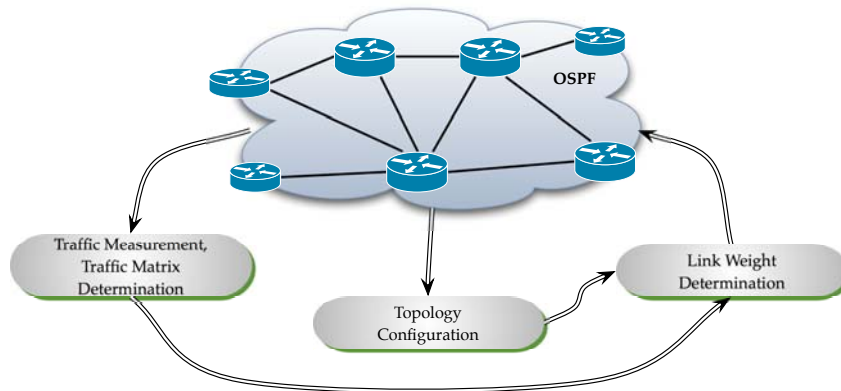


Fig. 5. Framework for IP traffic engineering.

It may be noted that for a moderately small network provider that deploys OSPF or IS-IS routing, link weight updating might not be a major issue. This is often the case when the provider has plenty of bandwidth, and some traffic fluctuation is tolerable. Thus, a hop-based metric is acceptable, or, alternately, a metric based on propagation delay or link speed can be used.

An alternate approach to IP network traffic engineering is to deploy multi-protocol label switching (MPLS) technology as the underlying technology in an IP/MPLS framework. An advantage of the MPLS approach is that it provides flexibility in controlling traffic, especially for different customer groups, based on service level agreements (SLAs). For example, MPLS tunnels may be set up with an appropriate bandwidth requirement to meet SLAs for different customers. Note that for optimal use of network resources, MPLS traffic engineering is necessary [24].

IP Prefix Management An important role of access ISPs is that they need to know and maintain IP prefixes for which they are the destination. For instance, they might obtain an address block from one of the Regional Internet registries (RIR), which they might slice and allocate to its customers. Alternately, a customer who already has an IP address block may want service from a particular provider. Either way, the IP prefixes are homed to a provider's AS. In earlier years, an edge provider (a provider that is at the "edge") depended on its upstream provider to forward any packet that was not destined for an IP address within its administrative domain. Thus, IP prefix maintenance was not an important issue. However, in recent years, due to IP address spoofing and to contain a packet from leaving a particular network destined for an invalid IP address, most edge providers now try to maintain the full IP prefix list of valid addresses. While they can receive IP prefixes through automated BGP advertisement, they may want to associate this information with other information about valid IP prefixes that they may have obtained through external mechanisms; this way, they can avoid sending packets to bogus IP prefixes. Thus, collection and maintenance of such external information is now an important management function most Internet service providers have (or need to have). Note that, for large tier-1 ISPs having a large number of customers with IP prefixes (which are directly connected), automated configuration management of customers is also an important management function [17].

Policy Management A third key issue is policy management. Most providers, whether they are tier-1 providers or transit providers, might or might not allow a specific IP prefix to use it as a bearer. Furthermore, for a specific IP prefix, it might provide preferential treatment regarding which upstream provider (identified by the AS number) is to be used. All in all, a provider is required to maintain two sets of policies: an import policy and an export policy. Based on business or other reasons, such policies might change over time. Thus, policy management is a major issue for many providers; in particular, this is quite complex for tier-1 providers. Illustrative examples of policies are listed in Table 3. For additional discussion on policy aspect, see [24].

4 PSTN and Internet Interworking: VoIP Routing

With the growth of the Internet, the need and interest to provide packetized voice service has grown in the past decade. With this development, the need has come for PSTN and Internet interworking so that voice over IP (VoIP) services can be provided for which the origination or termination of a call is possible either in the Internet or the PSTN. In this section, we briefly comment on management functionalities that may be needed due to routing interworking of the PSTN and the Internet (Figure 6).

In a PSTN-Internet interworking environment, call set up signaling information is to be communicated from one network to the other. For example, in the Internet, the session initiation protocol (SIP) is used for call set up messages. In the PSTN, ISDN User Part (ISUP) signalling is used for call set up messages using the SS7 signaling network. Once a call originating in the Internet arrives at the interworking gateway, a SIP message would need to be converted to ISUP signaling. Similarly, in a call originating in the PSTN that is destined for the Internet, the call signaling would need to be converted from an ISUP message to a SIP message. In addition to call signaling, there is also the media conversion step; that is, a call arriving from the PSTN side will likely use pulse-code modulation (PCM) coding; at the gateway, this would need to be converted to use over real-time transport protocol (RTP) with appropriate coding mechanism. From this discussion, it should be clear that the interworking gateway node is a critical component in this architecture. Thus, monitoring and configuration management of the gateway is critical for VoIP routing.

Secondly, as we discussed earlier, control schemes for the PSTN might also need to be handled for the Internet side. Furthermore, the IP side of the gateway is completely exposed to the Internet and can be a target for attacks such as denial of service attacks; thus, the gateway node would need to have the intrusion detection system for smooth functioning. Therefore, there are several emerging management functions that are appropriate for routing management of VoIP services.




Fig. 6. PSTN-Internet interworking.

5 Summary

In this paper, we presented the role of management in routing, primarily from a historical perspective. In particular, we have considered the PSTN and the Internet, and how the roles of management and management functions have evolved over time. There are the following common themes to note:

- When hierarchical routing was first introduced in the PSTN, it was not possible to do routing management since there were no communication functions to inform route availability information (this later changed with the availability of the signaling network). Initially, for the Internet also, it was not envisioned that routing management was necessary since the routing protocol had the adaptive mechanisms to address any changes automatically. From a network performance perspective, it was clear later that routing management and related functions are necessary.
- Dynamic call routing, when introduced in the PSTN, required a new set of management functionalities such as traffic monitoring, and measurements, and optimal route selection. Similarly, with further deployment experience with the Internet, it was recognized that large providers must do traffic-aware routing selection for optimal network performance; this required traffic monitoring, measurements, and determination of link weights.
- In both the PSTN and the Internet, address space-related management for efficient routing has become critical in recent years. For the PSTN, this arises due to the emergence of number portability, which allows a customer to change providers; this then requires ways to manage which address belongs to which provider and to find an efficient routing path. Similarly, in the Internet, an organization that is allocated an IP prefix may choose to change its network connectivity provider; this then results in announcement of route information and related management issues. It may be noted that for both the PSTN and the Internet, the multi-provider environment has necessitated the need for address-space related management.

As we can see, while certain functional aspects are network dependent, there are some common themes that are applicable to both networks.

The impact on routing due to emergence of multiple providers is important to note as well. In the case of the PSTN, the multi-provider environment first emerged in the long-distance segment of the network during the 1980s; different providers then implemented different dynamic routing schemes in the long-distance part. In recent years, the multi-provider setting has become common for network access at the end user level; along with number portability, the relation of addressing to routing has taken a new direction in the access network. In the case of the Internet, EGP was first deployed in the NSFNET, which served as the sole backbone. It then evolved to a multi-provider setting that was facilitated by BGP. This direction also has created interesting relations between addressing and routing, especially from a policy perspective. Thus, in a multi-provider setting, policy-based routing for both the PSTN and the Internet has emerged. Clearly, such evolution in routing has required new ways to address routing management.

A growing interest in both the PSTN and the Internet is the multi-domain cooperative management paradigm as no single provider has full control of the entire path for a user's traffic from origin to destination. Therefore, inter-domain traffic engineering and control for route management are becoming more and more important. In turn, this leads to control plane functionality playing a critical role in routing management. Thus, we envision seeing interesting challenges in years to come.

Acknowledgement

This paper relies on the book [24] with regard to the general discussion about routing. The figures are reproduced from the book, which Karthikeyan Ramasamy drew. Jennifer Rexford and Cory Beard read a draft of the paper and provided valuable feedback. I also thank the anonymous reviewers for their instructive comments and Nan Lorenz for carefully proofreading the paper.

References

1. G. R. Ash, *Dynamic Routing in Telecommunication Networks*. McGraw-Hill, 1997.
2. G. R. Ash, R. H. Cardwell, and R. P. Murray, "Design and optimization of networks with dynamic routing," *Bell System Technical Journal*, vol. 60, pp. 1787–1820, 1981.
3. G. R. Ash and P. Chemouil, "20 years of dynamic routing in telephone networks: Looking backward to the future," *IEEE Global Communications Newsletter*, pp. 1–4, October 2004, note: appears as insert in the October 2005 issue of *IEEE Communications Magazine*.
4. G. R. Ash, J. S. Chen, A. E. Frey, and B. D. Huang, "Real time network routing in a dynamic class-of-service network," in *Proc. 13th International Teletraffic Congress (ITC13)*, Copenhagen, Denmark, 1991, pp. 187–194.
5. H.-W. Braun, "Models of policy based routing," *IETF RFC 1104*, June 1989. <http://www.rfc-editor.org/rfc/rfc1104.txt>
6. H.-W. Braun, "The NSFNET routing architecture," *IETF RFC 1093*, February 1989. <http://www.rfc-editor.org/rfc/rfc1093.txt>
7. D. D. Clark, "Policy routing in Internet protocols," *IETF RFC 1102*, May 1989. <http://www.rfc-editor.org/rfc/rfc1102.txt>
8. J. Doyle and J. Carroll, *Routing TCP/IP, Volume II*. Cisco Press, 2001.
9. J. Doyle and J. Carroll, *Routing TCP/IP, Volume I, 2nd Edition*. Cisco Press, 2006.
10. A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "NetScope: Traffic engineering in IP networks," *IEEE Network*, vol. 14, no. 2, pp. 11–19, March/April 2000.
11. A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: Methodology and experience," *IEEE/ACM Transactions on Networking*, vol. 9, pp. 265–279, 2001, (an earlier version appeared in *Proc. ACM SIGCOMM'2000*).
12. B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *Proc. IEEE INFOCOM'2000*, Tel Aviv, Israel, March 2000, pp. 519–528.
13. M. Foster, T. McGarry, and J. Yu, "Number portability in the Global Switched Telephone Network (GSTN): An overview," *IETF RFC 3482*, February 2003. <http://www.rfc-editor.org/rfc/rfc3482.txt>

14. V. Fuller, T. Li, J. Yu, and K. Varadhan, "Supernetting: An address assignment and aggregation strategy," *IETF RFC 1338*, June 1992. <http://www.rfc-editor.org/rfc/rfc1338.txt>
15. R. J. Gibbens, F. P. Kelly, and P. B. Key, "Dynamic Alternate Routing—modeling and behaviour," in *Proc. 12th International Teletraffic Congress (ITC12)*, Turin, Italy, 1988, pp. 3.4A3.1–3.4A3.7.
16. A. Girard, *Routing and Dimensioning in Circuit-Switched Networks*. Addison-Wesley, 1990.
17. J. Gottlieb, A. Greenberg, J. Rexford, and J. Wang, "Automated provisioning of BGP customers," *IEEE Network*, vol. 17, no. 6, pp. 44–55, November/December 2003.
18. D. Haenschke, D. A. Kettler, and E. Oberer, "Network management and congestion in the U.S. telecommunications network," *IEEE Trans. on Communications*, vol. COM-29, pp. 376–385, 1981.
19. C. Huitema, *Routing in the Internet, 2nd Edition*. Prentice-Hall, 2000.
20. B. Leiner, "Policy issues in interconnecting networks," *IETF RFC 1124*, September 1989, (available only as postscript or PDF file). <http://www.faqs.org/rfc/rfc1124.pdf>
21. B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff, "A brief history of the Internet," version 3.32, last revised 10 December, 2003. <http://www.isoc.org/internet/history/brief.shtml>
22. K. Lougheed and Y. Rekhter, "A Border Gateway Protocol (BGP)," *IETF RFC 1105*, June 1989. <http://www.rfc-editor.org/rfc/rfc1105.txt>
23. K. Lougheed and Y. Rekhter, "A Border Gateway Protocol 3 (BGP-3)," *IETF RFC 1267*, October 1991. <http://www.rfc-editor.org/rfc/rfc1267.txt>
24. D. Medhi and K. Ramasamy, *Network Routing: Algorithms, Protocols, and Architectures*. Morgan Kaufmann Publishers (an imprint of Elsevier), 2007.
25. D. L. Mills, "Exterior gateway protocol formal specification," *IETF RFC 904*, April 1984. <http://www.rfc-editor.org/rfc/rfc904.txt>
26. D. L. Mills, "Autonomous confederations," *IETF RFC 975*, February 1986. <http://www.rfc-editor.org/rfc/rfc975.txt>
27. M. Pióro, A. Szentesi, J. Harmatos, A. Jüttner, P. Gajowniczek, and S. Kozdrowski, "On open shortest path first related network optimization problems," in *Proc. IFIP ATM IP 2000*, Ilkley, England, July 2000, (see also [28]).
28. M. Pióro, A. Szentesi, J. Harmatos, A. Jüttner, P. Gajowniczek, and S. Kozdrowski, "On open shortest path first related network optimization problems," *Performance Evaluation*, vol. 48, pp. 201–223, 2002, (see [27] for a preliminary version).
29. D. M. Piscitello and A. L. Chapin, *Open Systems Networking: TCP/IP and OSI*. Addison-Wesley, 1993.
30. Y. Rekhter, "EGP and policy based routing in the new NSFNET backbone," *IETF RFC 1092*, February 1989. <http://www.rfc-editor.org/rfc/rfc1092.txt>
31. Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," *IETF RFC 1771*, March 1995, (Made obsolete by [32]). <http://www.rfc-editor.org/rfc/rfc1771.txt>
32. Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," *IETF RFC 4271*, January 2006. <http://www.rfc-editor.org/rfc/rfc4271.txt>
33. J. Rexford, "Route optimization in IP networks," in *Handbook of Optimization in Telecommunications*. Springer, 2006, pp. 679–700, M. G. C. Resende and P. Pardalos (Eds.).
34. R. F. Rey (Ed.), *Engineering and Operations in the Bell System, 2nd Edition*. AT&T Bell Laboratories, 1983.
35. E. Rosen, "Exterior gateway protocol (EGP)," *IETF RFC 827*, October 1982. <http://www.rfc-editor.org/rfc/rfc827.txt>

36. L. J. Seamonson and E. Rosen, ““Stub” exterior gateway protocol,” *IETF RFC 888*, January 1984. <http://www.rfc-editor.org/rfc/rfc888.txt>
37. Syniverse Technologies, “A global perspective on number portability,” May 2004. <http://www.syniverse.com/pdfs/MNPRreport.pdf>
38. E. Szybicki and M. Lavigne, “Alternate routing for a telephone system,” U.S. Patent No. 4,284,852, August 18, 1981.
39. Telecom Regulatory Authority of India, “Recommendation on mobile number portability,” March 2006. <http://www.trai.gov.in/recomm8mar06.pdf>
40. D. M. Tow, “Network management—recent advances and future trends,” *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 4, pp. 732–741, 1988.
41. J. H. Weber, “A simulation study of routing and control in communication networks,” *Bell Systems Technical Journal*, vol. 43, pp. 2639–2676, 1964.