Chapter 8

Routing Symmetry

We now analyze the routes from our measurement study to assess the degree to which routes are *symmetric*. We first motivate the investigation by discussing the impact of routing asymmetry on different network protocols and measurements. We then give an overview of various mechanisms that can introduce asymmetry into Internet routing, including “hot potato” routing (§ 8.2), which could result in a greater proportion of asymmetric routes in the future. We next introduce a definition of routing symmetry, and show that practical considerations require a revision in which we view routes as asymmetric only if they visit different cities or autonomous systems. We then assess our data for these asymmetries and find that, overall, 50% of the time an Internet path includes a major asymmetry in terms of the cities visited in the different directions, and 30% of the time it includes a major asymmetry in terms of autonomous systems visited. We finish with a discussion of the magnitude of the asymmetries, most of which differ at just one “hop,” but some at many hops.

8.1 Importance of routing symmetry

Routing symmetry affects a number of aspects of network behavior. When attempting to assess the one-way propagation time between two Internet hosts, the common practice is to assume it is well approximated as half of the round-trip time (RTT) between the hosts [CPB93a]. The Network Time Protocol (NTP) needs to make such an assumption when synchronizing clocks between widely separated hosts [Mi92a]. If routes are asymmetric, however, the assumption might easily lead to error. The NTP design utilizes multiple time server peers and robust algorithms to choose among them for the best time offset to use to account for propagation effects. Thus, routing asymmetry has an impact on NTP only if the paths between two NTP communities are predominantly asymmetric, with similar differences in one-way times. In that case, the two communities will keep consistent time among themselves, but not between each other.[1]

Claffy and colleagues studied variations in one-way latencies between the United States, Europe, and Japan [CPB93a]. They discuss the difficulties of measuring *absolute* differences in propagation times in the absence of separately-synchronized clocks, but for their study they focussed on *variations*, which does not require synchronization of the clocks. They found that the

---

[1] Recently, however, highly accurate atomic clocks have become much more affordable than in the past (as have Global Positioning System receivers, which also provide reliable time). These provide an independent solution to the problem of keeping widely separated NTP servers synchronized.

two opposing directions of a path do indeed exhibit considerably different latencies, in part due to different congestion levels, and in part due to routing changes, which they detected using the TTL method (§ 7.7).

Along with affecting Internet protocols such as NTP, routing asymmetry can render network measurement considerably more difficult. Often it is easiest to perform measurements at a single endpoint of a network path, but in the face of routing asymmetries, such measurements might be unable to distinguish between considerably different behavior along the forward and reverse directions of the path. We explore this problem at length in Part II (see § 9.1.3 for a general discussion).

Closely related to this measurement problem, routing asymmetry also potentially complicates mechanisms by which connection endpoints can infer network conditions from the pattern of packet arrivals they observe. For example, we develop a technique in Chapter 14 for estimating the "bottleneck bandwidth" of the network path used by a connection. The technique works by examining the timing with which packets arrive at their receiver. If routing is symmetric, then (for most link technologies) the bottleneck bandwidth measured by this technique will be the same as that encountered by packets sent in the other direction. Symmetry could, for example, allow the server for a request/reply application such as the World Wide Web [BCLF+], or, more generally, T/TCP [Br94], to determine the link bandwidth available for sending its reply, based on the bandwidth inferred from the request. If routing is asymmetric, however, then the server runs the risk of inferring an incorrect value for the bandwidth.[2] However, we show in Chapters 14 and 16 that bottleneck bandwidths and delays are often asymmetric along the two directions of a path, and attribute the difference at least in part to routing asymmetries.

Finally, recent work has investigated the characteristics of network traffic *flows* as viewed by a router [CBP95]. That study describes a taxonomy of methodologies that can be used by routers to define and manage flow state. One finding of the study is that a large number of flows are bidirectional, due in part to request/reply transactions such as those used by the Domain Name System (DNS; [MD88]) and the World Wide Web. When a router R sees a flow likely to be bidirectional, for example a DNS request from A to B, one might consider establishing *anticipatory flow state* in the router for the reply coming from B to A, to avoid the overhead of two separate trips through the "slow path" associated with flows for which there is no cached state. With prevalent routing asymmetry, however, while B may very likely send such a message shortly, the reply could well *not* be routed via R, in which case the anticipatory flow state is wasted effort and resources.

Similarly, *accounting* used to charge for carrying network traffic is complicated by the possibility of locally observing only one direction of a traffic flow. For example, a recently developed architecture for Internet traffic flow measurement has a basic assumption that routers observe bidirectional flows [BMR97].

8.2 Sources of routing asymmetries

In this section we discuss several mechanisms that can lead to routing asymmetries. To illustrate, we assume the viewpoint of a router R_0 faced with the decision of how to forward packets originated by host A and destined for host B. In addition to the upstream router from which R_0

---

[2] Even if routing is symmetric, the server cannot rely on the congestion levels being symmetric. Thus, as with routing stability, routing symmetry is *necessary* but not *sufficient* for predicting network behavior.

receives packets sent by A, R_0 is connected to two potential downstream routers, R_1 and R_2, and the decision it must make is to which of these it forwards packets bound for B. Let us also assume that packets from B headed to A arrive at R_0 via R_1 (but in general R_0 does not itself know this fact), and that these packets first pass through a router R_3, which makes the decision whether to use the route that ultimately delivers the packets to R_0 via R_1, or a different route that results in the packets arriving at R_0 via R_2.

In general, routing algorithms incorporate "link costs" or *metrics* to quantify the desirability of using a particular link for a given route [Pe92, St95]. To assure reliable operation, a router also generally knows of multiple paths available to a remote destination B, so we assume that R_0 has two metrics, μ_1 and μ_2, associated with forwarding packets to B via R_1 or R_2. If $\mu_1 = \mu_2$, then R_0 must somehow arbitrate between them. If it does so deterministically, by picking R_2, then an asymmetry is created.[3]

Another way of introducing asymmetry is via configuration asymmetries or errors. For example, if due to misconfiguration R_0 believes that using the link to R_1 is very expensive, but R_1 does not share this view, then R_0 will artificially inflate the cost of using R_1 to get to B, and instead pick R_2.

Network topology changes can also introduce routing asymmetries, albeit transient ones, due to the non-negligible amount of time required for changes to propagate through the network. For example, suppose R_2 learns of a better route to R_3 than it had before. If knowledge of this new route propagates to R_0 before R_3, then R_0 will switch from R_1 to R_2, and an asymmetry will exist until R_3 learns of the route.

Another transient mechanism for creating routing asymmetries can arise due to *adaptive routing* (§ 7.2), in which a router attempts to shift traffic from a highly loaded link to a less loaded link. For example, R_0 might decide that it is sending too much traffic via the link to R_1 (the bulk of this traffic might not be destined for B), so it increases the metrics associated with R_1 to the point where routing via R_2 becomes the preferred route to B. More generally, if routing metrics include a notion of current congestion levels, then asymmetric congestion in the network can lead to asymmetric routing, as the network alters its routing to avoid the congested region.

A final mechanism introducing asymmetry, and one of possibly growing importance, concerns "hot potato" and "cold potato" routing. In the past, Internet backbones were primarily operated by a single entity. In recent years this has changed, with the growth of competing Internet Service Providers (ISP's) due to the privatization of the Internet infrastructure.

Suppose host A in California uses ISP I_A, and host B in New York uses I_B. Assume that both I_A and I_B provide Internet connectivity across the entire United States. When A sends a packet to B, the routers belonging to I_A must at some point transfer the packet to routers belonging to I_B. Since cross-country links are a scarce resource, both I_A and I_B would prefer that the other convey the packet across the country. If the inter-ISP routing scheme allows the upstream ISP (I_A, in our example) to determine when to transfer the packet to I_B, then, due to the preference of avoiding the cross-country haul, I_A will elect to route the packet via I_B as soon as possible. This form of routing is known as "hot potato." In our example, it leads to I_A transferring the packet to I_B in California. But when B sends traffic to A, I_B gets to make the decision as to when to forward the traffic to I_A, and with hot potato it will choose to do so in New York. Since the paths between California and New York used by I_A and I_B will in general be quite different, hot potato routing thus leads to a

---

[3] If it alternates between R_1 and R_2, it creates *fluttering*, as discussed in § 6.6.

major routing asymmetry between A and B.

Conversely, if the *downstream* ISP can control where the upstream ISP transfers packets to it, then the result is “cold potato” routing, in which I_B instructs I_A that, to reach B, I_A should forward packets to I_B's New York network access point (NAP). Similarly, I_A advertises to I_B that, to reach A, I_B should forward packets to I_A's California NAP. The result is that packets from A to B travel across the country via I_A's links, while those from B to A travel via I_B's links. The paths are the opposite of those resulting from hot potato routing, but the degree of asymmetry remains the same, and potentially large.

For further discussion of asymmetry issues, see [Che95].

8.3 Definition of routing symmetry

In this section we develop a definition for whether two routes are symmetric. We first try the following:

Definition 1 *For two hosts A and B, let $r_1, \ldots, r_n$ denote the routers visited in sequence by packets sent from A to B, and $r'_1, \ldots, r'_m$ denote those visited in sequence by packets from B to A. Then the two routes are* symmetric *if and only if $n = m$ and:*

$$\forall i, 1 \leq i \leq n : r_i = r'_{n+1-i}.$$

Definition 1 presents two problems. First, for routes considered asymmetric, the definition fails to provide a notion of the *degree* of asymmetry. For example, if a site has two Internet access points, then we could find that traffic from A to B leaves the site at the first access point for a downstream router R, while traffic from B to A comes to the site also from R, but arriving at the second access point. Such an asymmetry is minor. For example, it will have minimal impact on the accuracy of the NTP protocol (§ 8.1). On the other hand, if the route from A to B visits a different *city* than does the route from B to A, then the two paths might have considerably different properties, and the asymmetry is major.

To illustrate these differences, consider the route we observed in $\mathcal{R}_1$ from `ucol` to `ucl` (where we have annotated the cities visited in parentheses), shown in Figure 8.1. One of the complementary routes we observed from `ucl` to `ucol` is shown in Figure 8.2. This route visits the same cities as the reverse route, though not the same routers; the asymmetry is minor. On the other hand, we also observed a route from `ucl` to `ucol` as shown in Figure 8.3. In this case, the detour via California is skipped, shaving perhaps 2,000 kilometers of travel from the route: a major asymmetry.

A second problem with Definition 1 is determining whether two routers r_i and r'_j are indeed the same router. The difficulty arises because `traceroute` provides an IP address for each hop, but these do not uniquely identify routers. In general, routers have multiple IP addresses, one for each network interface attached to the router. Furthermore, these IP addresses can translate to different hostnames. Thus, for example, it is difficult to determine whether the IP address with hostname `sl-ana-3-s2/4-t1.sprintlink.net` in Figure 8.1 corresponds to the same router as that with hostname `sl-ana-3-f0/0.sprintlink.net` in Figure 8.2.

We address both these difficulties using a revised definition:

```
cs-gw-discovery.cs.colorado.edu  (Boulder, CO)
cu-gw.colorado.edu
sl-ana-3-s2/4-t1.sprintlink.net  (Anaheim, CA)
sl-ana-1-f0/0.sprintlink.net
sl-fw-6-h2/0-t3.sprintlink.net   (Fort Worth, TX)
sl-fw-5-f1/0.sprintlink.net
sl-dc-8-h3/0-t3.sprintlink.net   (Washington, D.C.)
icm-dc-1-f0/0.icp.net
icm-london-1-s1-1984k.icp.net    (London, UK)
smds-gw.ulcc.ja.net
smds-gw.ucl.ja.net
cisco-pb.ucl.ac.uk
cisco.cs.ucl.ac.uk
neptune.cs.ucl.ac.uk
```

Figure 8.1: Route observed from `ucol` to `ucl`

```
cisco.cs.ucl.ac.uk               (London, UK)
cisco-pb.ucl.ac.uk
cisco-b.ucl.ac.uk
gw.lon.ja.net
eu-gw.ja.net
icm-lon-1.icp.net
icm-dc-1-s3/2-1984k.icp.net      (Washington, D.C.)
sl-dc-6-f0/0.sprintlink.net
sl-dc-8-f0/0.sprintlink.net
sl-fw-5-h4/0-t3.sprintlink.net   (Fort Worth, TX)
sl-fw-6-f0/0.sprintlink.net
sl-ana-1-h2/0-t3.sprintlink.net  (Anaheim, CA)
sl-ana-3-f0/0.sprintlink.net
sl-ucb-2-s0-t1.sprintlink.net    (Boulder, CO)
cs-gw.colorado.edu
clark.cs.colorado.edu
```

Figure 8.2: Route observed from `ucl` to `ucol`

```
cisco.cs.ucl.ac.uk                 (London, UK)
cisco-pb.ucl.ac.uk
cisco-c.ucl.ac.uk
smds-gw.ulcc.ja.net
icm-lon-1.icp.net
icm-dc-1-s3/2-1984k.icp.net        (Washington, D.C.)
sl-dc-8-f0/0.sprintlink.net
sl-fw-5-h4/0-t3.sprintlink.net     (Fort Worth, TX)
sl-fw-4-f0/0.sprintlink.net
sl-ucb-1-s0-t1.sprintlink.net      (Boulder, CO)
cns-gw-suns.colorado.edu
cs-gw.colorado.edu
lewis.cs.colorado.edu
```

Figure 8.3: Second route observed from `ucl` to `ucol`

Definition 2 *For two hosts A and B, let $c_1, \ldots, c_n$ denote the* cities *visited in sequence by packets sent from A to B, and $c'_1, \ldots, c'_m$ denote those visited in sequence by packets from B to A. Then the two routes are* symmetric *if and only if $n = m$ and:*

$$\forall i, 1 \leq i \leq n : c_i = c'_{n+1-i}.$$

This definition deals with the first difficulty of the original definition by discarding all minor routing asymmetries—we consider a routing asymmetry interesting only if it is major. It resolves the second difficulty because it is considerably easier to tell whether two IP addresses are located in the same city than whether they refer to the same router, since with a bit of effort it is generally possible to determine the city corresponding to an Internet hostname (cf. § 5.3). For example, we know from the Sprintlink naming convention that both `sl-ana-3-s2/4-t1.sprintlink.net` and `sl-ana-3-f0/0.sprintlink.net` are located in Anaheim, California.

We can make an analogous definition for routes differing in the autonomous systems they visit, rather than the cities.

8.4 Analysis of routing symmetry

In $\mathcal{R}_1$, we did not make simultaneous measurements of the paths $A \Rightarrow B$ and $B \Rightarrow A$, which introduces ambiguity into an analysis of routing symmetry: if a measurement of $A \Rightarrow B$ is asymmetric to a later measurement of $B \Rightarrow A$, is that because the route is the same but asymmetric, or because the route changed?

In $\mathcal{R}_2$, however, the bulk of the measurements were *paired*: we first measured $A \Rightarrow B$ and then immediately afterward measured $B \Rightarrow A$. Barring rapid route oscillations (which we can avoid by eliminating pathological `traceroutes` from our analysis), these measurements allow us to unambiguously determine whether the route between A and B is symmetric.

The $\mathcal{R}_2$ measurements contain 11,339 successful pairs of measurements, in which we were able to conduct `traceroutes` in both directions between sites A and B, neither of the measurements encountering pathologies.

We find that *49% of the measurements observed an asymmetric path that visited at least one different city*.

There is a large range, however, in the prevalence of asymmetric routes among paths to and from the different sites. For example, 86% of the paths involving `umann` were asymmetric, because nearly all outbound traffic from `umann` travel via Heidelberg, but none of the inbound traffic does. At the other end of the spectrum, only 25% of the paths involving `umont` were asymmetric (but this is still a significant amount).

If we consider autonomous systems rather than cities, then we still find asymmetry quite common: about 30% of the paired measurements observed different autonomous systems traversed in the path's two directions. The most common asymmetry was the addition of a single AS in one of the directions. This can reflect a major change, however. For example, the most common of these additions was the presence of SprintLink routers in one direction along the path but not in the other.

Again, we find a wide range in the prevalence of asymmetry among the different sites. Fully 84% of the paths involving the `ucl` site were asymmetric, mostly due to some paths including JANET routers in London and others not (unsurprising, given the rapid oscillation between JANET and non-JANET routers discussed in § 7.6.1). On the other end of the spectrum, only 7.5% of `adv`'s paths were asymmetric at AS granularity.

8.5 Increasing prevalence of asymmetry

We previously analyzed $\mathcal{R}_1$ for routing asymmetry, attempting to adjust for the non-simultaneity of its measurements by only using measurements spaced less than a day apart. The mismatch is likely to overestimate routing asymmetry, since if the route changes between measurements that may be incorrectly regarded as an asymmetry, per our discussion at the beginning of § 8.4. The mismatch can also introduce false symmetries, if the route happens to change to the symmetric counterpart, but this circumstance is probably more rare than introducing false asymmetries.

In the $\mathcal{R}_1$ measurements, we found 30% of the paths contained city-level asymmetries. The large discrepancy between this figure and the 50% figure for the $\mathcal{R}_2$ measurements suggests that over the course of a year routing became significantly more asymmetric. We surmise that the increase of asymmetry is likely due to the “hot potato” effect discussed in § 8.2. If so, then the rise in asymmetry has its roots in commercial factors, and frequent routing asymmetry may continue to be common in the Internet in the future. From a measurement perspective, this would be unfortunate, for the reasons given § 8.1, and further developed in § 9.1.3.

8.6 Size of asymmetries

We finish our study of routing symmetry with a look at the size of the different asymmetries. We can assign a “magnitude” to each asymmetry in terms of the number of cities different in the two directions. We consider each “city hop” at which the two directions of a path differ as contributing a magnitude of 1; if one direction has more “city hops” than the other, each additional city contributes $\frac{1}{2}$. For example, for the paths between `rain` and `bnl`, we observed simultaneous measurements of the following routes:

```
r0.pdx.rain.rg.net                 (Portland)
```

```
sl-stk-13-s2/2-t1.sprintlink.net  (Stockton)
sl-stk-5-f0/0.sprintlink.net
sl-dc-6-h1/0-t3.sprintlink.net    (Washington, D.C)
sl-pen-1-h2/0-t3.sprintlink.net   (Pennsauken)
sl-pen-2-f0/0.sprintlink.net
ny-nyc-2-h1/0-t3.nysernet.net     (New York)
ny-nyc-6-f0/0.nysernet.net
ny-dp-1-h0/0-t3.nysernet.net      (Deer Park)
ny-bnl-2-s0-t1.nysernet.net       (BNL)
cerberus.bnl.gov
frog.rhic.bnl.gov
```

and

```
cerberus.90.bnl.gov               (BNL)
nioh.bnl.gov
192.12.15.224
llnl-satm.es.net                  (Livermore)
ames-llnl.es.net                  (Mountain View)
fix-west-cpe.sanfrancisco.mci.net (San Francisco)
borderx2-hssi2-0.sanfrancisco.mci.net
core2-fddi-1.sanfrancisco.mci.net
core1-hssi-2.sacramento.mci.net   (Sacramento)
core-hssi-3.seattle.mci.net       (Seattle)
border1-fddi-0.seattle.mci.net
rgnet-b1-serial2-3.seattle.mci.net
chia.rain.net                     (Portland)
```

The paths differ at five “city hops,” Stockton/Seattle, Washington/Sacramento, Pennsauken/San Francisco, New York/Mountain View, and Deer Park/Livermore, so we assign a magnitude of 5 to this asymmetry.

Figure 8.4 shows the distribution of asymmetry magnitudes. We see that the asymmetries typically include only one different city hop, or, even more commonly, just one additional city. About one third of the asymmetries have magnitude 2 or greater. We should bear in mind, though, that this corresponds to almost 20% of all the paired measurements in our study, and can correspond to a very large asymmetry. For example, a magnitude 2 asymmetry between `ucl` and `umann` differs at the central city hops of Amsterdam and Heidelberg in one direction, and Princeton and College Park in the other!

In general, the presence of such asymmetries highlights the difficulties of providing a consistent topological view in an environment as large and diverse as the Internet.

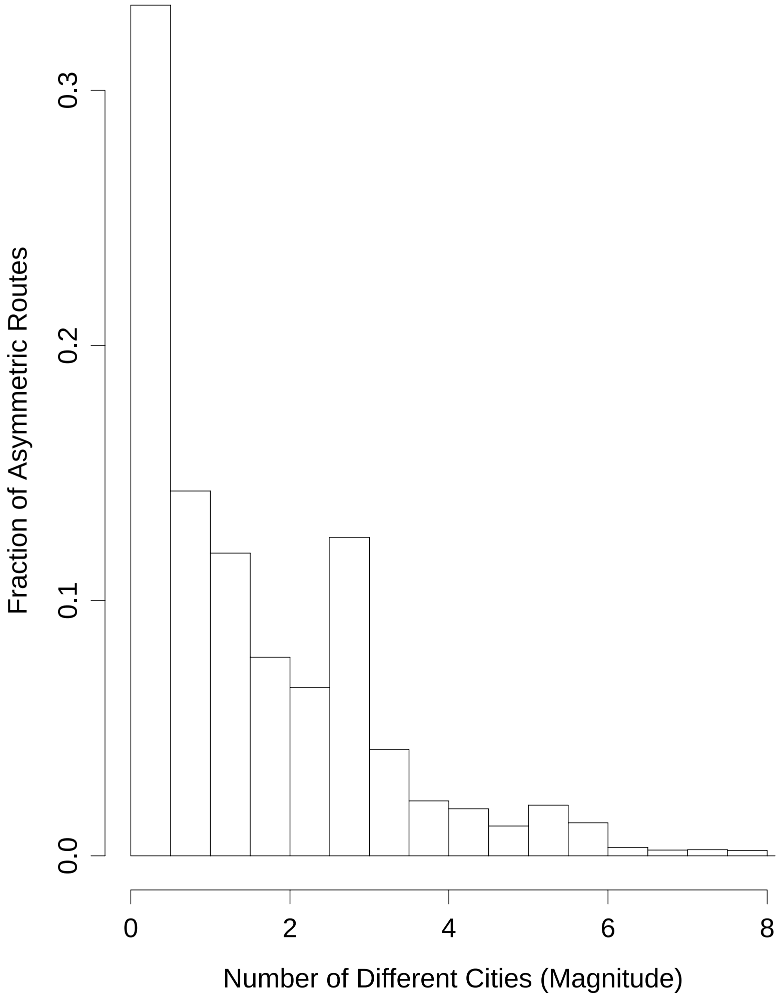

Figure 8.4: Distribution of asymmetry sizes