Chapter 5

The Raw Routing Data

In this chapter we discuss the sites that participated in our routing experiments, the duration of the experiments, and the preliminary reduction of the raw data we gathered.

5.1 Participating sites

The first routing experiment began the evening of Tuesday, November 8, 1994, and lasted until the morning of Saturday, December 24. During this time, we attempted 6,991 `traceroutes` between 27 sites. We refer to this collection of measurements as $\mathcal{R}_1$ (dataset #1). We will often refer to a single such measurement as a "`traceroute`."

The second experiment began the morning of Friday, November 3, 1995, and lasted until the afternoon of Thursday, December 21. It included 37,097 attempted `traceroutes` between 33 sites. We refer to this collection of measurements as $\mathcal{R}_2$. Details of the measurements and the sampling intervals are discussed in § 4.1. Both $\mathcal{R}_1$ and $\mathcal{R}_2$ are publicly available from the *Internet Traffic Archive*, at:

`http://www.acm.org/sigcomm/ITA`

under the name *NPD-Routes*.[1]

Table I lists the sites participating in $\mathcal{R}_1$, giving the abbreviation we will use to refer to each site, the site's Internet domain, the number of days it participated in the study, a brief description of the site, and its location. These sites also participated in $\mathcal{R}_2$, except for `batman`, `korea`, `usc`, and `xor`. Table II lists the additional sites participating in $\mathcal{R}_2$. In $\mathcal{R}_2$, all sites participated at least a month, except for `ukc`, which participated for 23 days, and 13 of the sites participated for the maximum of 48 days.

The sites include educational institutes, research labs, network service providers, and commercial companies, in 9 countries. Figures 5.1 and 5.2 show the geographic locations of the North American and European sites.

[1] At the time of this writing, the Archive is moving from its old location to the above URL. If the reader has any difficulty accessing the Archive, send email to `vern@ee.lbl.gov`.

Name	Domain	Days	Description	Location
austr	mu.oz.au	24	University of Melbourne	Melbourne, Australia
batman	batman.net	11	Experimental ATM network at National Center for Atmospheric Research	Boulder, CO
bnl	bnl.gov	37	Brookhaven National Lab	Brookhaven, NY
bsdi	bsdi.com	9	Berkeley Software Design, Inc.	Colorado Springs, CO
connix	connix.com	22	Caravela Software	Middlefield, CT
harv	harvard.edu	9	Harvard University	Cambridge, MA
inria	inria.fr	9	INRIA	Sophia, France
korea	postech.ac.kr	36	Pohang Institute of Science and Technology	Pohang, South Korea
lbl lbli	lbl.gov lbl.gov	45 45	Lawrence Berkeley Lab LBL home computer connected via ISDN	Berkeley, CA Berkeley, CA
mit	mit.edu	21	Massachusetts Institute of Technology	Cambridge, MA
ncar	ucar.edu	22	National Center for Atmospheric Research	Boulder, CO
nrao	cv.nrao.edu	44	National Radio Astronomy Observatory	Charlottesville, VA
oce	oce.nl	19	Oce-van der Grinten	Venlo, The Netherlands
pubnix	va.pubnix.com	11	Pix Technologies Corp.	Fairfax, VA
sdsc	sdsc.edu	24	San Diego Supercomputer Center	San Diego, CA
sri	sri.com	9	SRI International	Menlo Park, CA
ucl	ucl.ac.uk	24	University College	London, U.K.
ucol	colorado.edu	45	University of Colorado	Boulder, CO
ukc	ukc.ac.uk	24	University of Kent	Canterbury, U.K.
umann	uni-mannheim.de	19	University of Mannheim	Mannheim, Germany
umont	umontreal.ca	15	University of Montreal	Montreal, Canada
unij	kun.nl	9	University of Nijmegen	Nijmegen, The Netherlands
usc	usc.edu	45	University of Southern California	Los Angeles, CA
ustutt	uni-stuttgart.de	16	University of Stuttgart	Stuttgart, Germany
wustl	wustl.edu	33	Washington University	St. Louis, MO
xor	xor.com	30	XOR Network Engineering	East Boulder, CO

Table I: Sites participating in first experiment ($\mathcal{R}_1$)

Name	Domain	Description	Location
`adv`	advanced.org	Advanced Network & Services	Armonk, New York
`austr2`	newcastle.edu.au	University of Newcastle	Newcastle, Australia
`mid`	mid.net	MIDnet	Lincoln, Nebraska
`near`	near.net	NEARnet	Cambridge, Massachusetts
`panix`	nyc.access.net	Public Access Networks Corporation	New York, New York
`rain`	rain.net	RAINet, Inc.	Portland, Oregon
`sandia`	ca.sandia.gov	Sandia National Laboratories	Livermore, California
`sintef1`	sintef.no	University of Trondheim	Trondheim, Norway
`sintef2`	sintef.no	University of Trondheim	Trondheim, Norway
`ucla`	ucla.edu	University of California	Los Angeles, California

Table II: Additional sites participating in second experiment ($\mathcal{R}_2$)

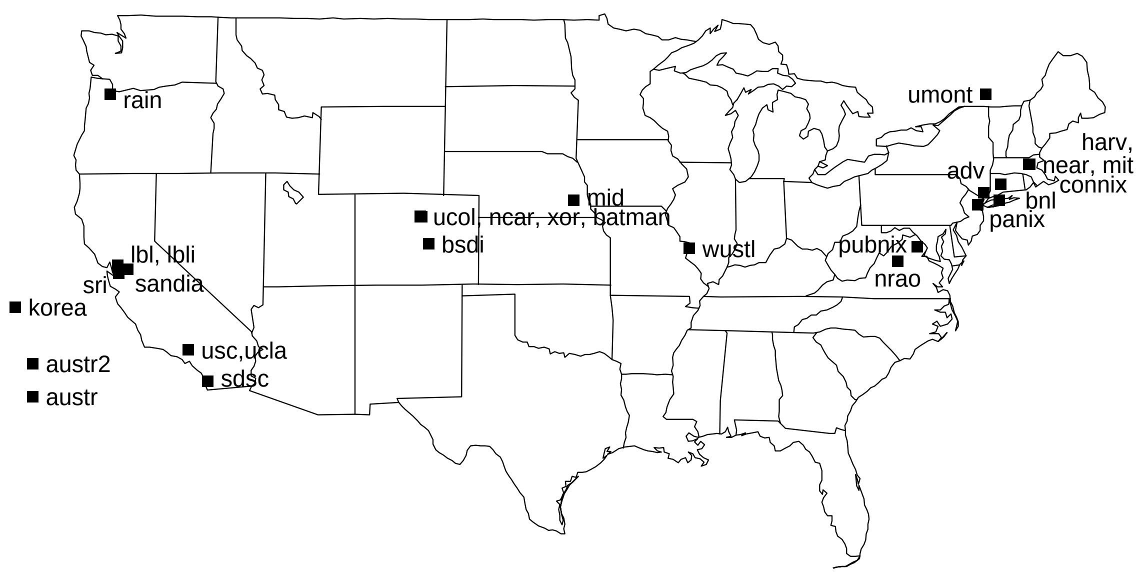

Figure 5.1: Sites participating in routing study, North America and Asia

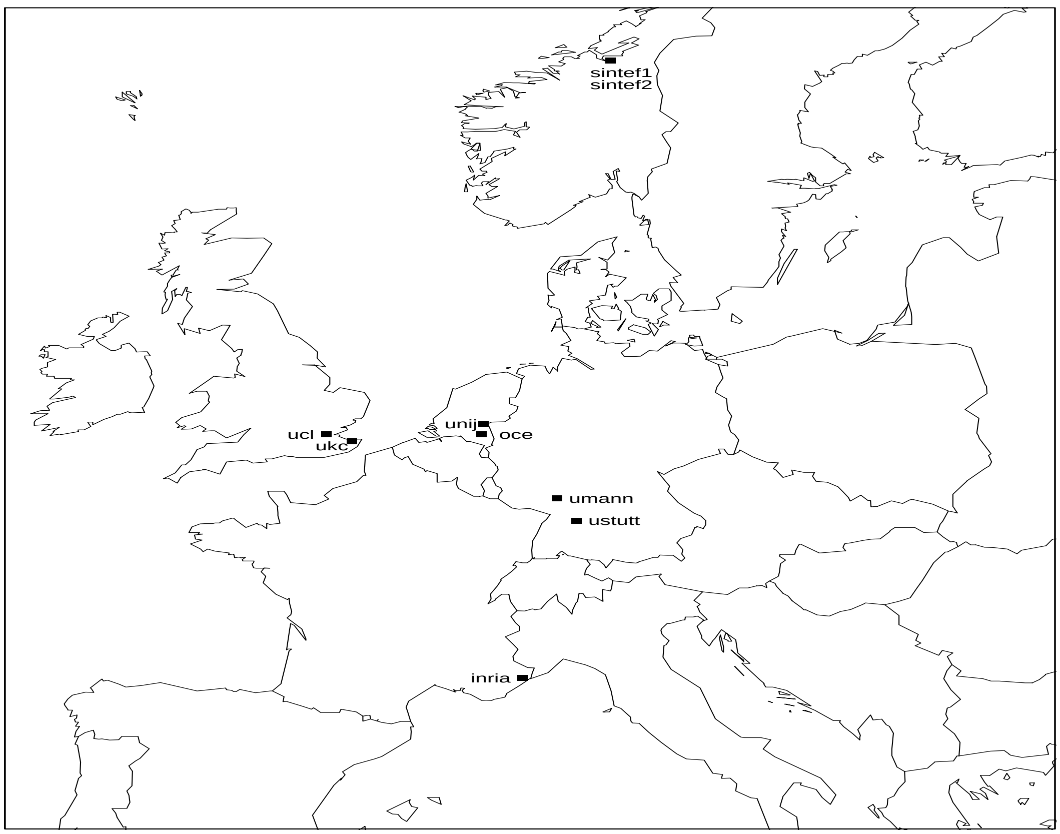

Figure 5.2: Sites participating in routing study, Europe

Status	Experiment 1 #	Experiment 1 %	Experiment 2 #	Experiment 2 %
Unable to contact daemon	495	7.1%	1,872	5.0%
Daemon configuration error	25	0.4%	15	0.04%
Host lookup failure	12	0.2%	101	0.3%
Total failures	532	7.6%	1,988	5.4%
Total successes	6,459	92.4%	35,109	94.6%
Total	6,991	100.0%	37,097	100.0%

Table III: Summary of routing experiment difficulties

5.2 Measurement breakdown

In the two experiments, between 5–8% of the `traceroutes` failed outright (i.e., we were unable to contact the remote `npd`, execute `traceroute` and retrieve its output). As shown in Table III, almost all of the failures were due to an inability of the `npd_control` process to contact the remote daemon. Some of these were failures involving `lbli`; that site, due to its ISDN link frequently being down (§ 6.7.4), was often unreachable. But for most of the failures we do not *a priori* know whether they represent the remote host being down or an Internet connectivity failure. It is important to note that, if the latter was frequently the case, then to some degree *the assumptions behind PASTA are invalid*, since an agent at the remote site with knowledge of current connectivity problems could reliably predict no sampling would occur in the near future (§ 4.3).

For our analysis, the effect of these failures to contact the remote daemon (`npd`) will lead to a bias towards *underestimating* Internet connectivity failures, because sometimes the failure to contact the remote daemon will result in losing an opportunity for a `traceroute` experiment to reveal the lack of connectivity between that site and another remote site that shares the same path as used between `npd_control` and the daemon.

When taking the $\mathcal{R}_2$ measurements, however, we somewhat corrected for this underestimation by *pairing* each measurement of the path $A \Rightarrow B$ with a measurement of the path $B \Rightarrow A$.[2] If `npd_control` was unable to reach one of either A or B, it still attempted to contact the other to measure the reverse route. In those circumstances where it was able to measure the reverse route, it still had an opportunity to observe the routing fault, if present in both directions.

`npd_control` was unable to reach one of either A or B 1,872 times. It was unable to contact the other host of the measurement pair, either, in only 5% of these instances. Thus, for the most part, the $\mathcal{R}_2$ measurements do not suffer from bias in observing bidirectional routing faults.

We could further reduce this measurement problem by introducing a "batch" design to `npd`, where the daemon would accept a list of measurements it should make at future points in time, and would email back the results when they were complete. We did not adopt this approach because one of our goals in the design of `npd` was to keep it simple enough that sites volunteering to run it could with reasonable ease inspect the code to see what they were running.

[2] About 20% of the measurements were not paired, because they were made in conjunction with the measurements discussed in Part II.

Destination \ Source	total	lbl	usc	nrao	ucol	lbli	ucl	bnl	wustl	xor	ukc	sdsc	mit	connix	ncar	oce	umont	umann	ustutt	harv	korea	batman	inria	sri	bsdi	unij	pubnix	austr
austr	100	4	6	4	9	1	4	3	8	3	9	5	3	5	2	6	4	2	5	2	3	1	1	4	1	3	2	
pubnix	98	3	4	5	2	5	6	5	1	5	6	2	3	4	1	8	1	4	4	5		3	4	4	5	7		1
unij	82	10	4		3	1	2	2	2	4	6	1	6	1	6	2	2	5	1	6		1	3	5	5		3	1
bsdi	87	3	3	3	3	4		5	6	4	4	1	1	1	1	1	7	3	5	6	2	7	5	5		3	3	1
sri	104	8	2	5	3	3	1	4	7	5	5	4	6	6	2	4	4	8	2	5		1	4		6	3	5	1
inria	106	7	6	4	6	1	4	3	5	6	4	3	4	4	4	4	2	6	4	8	1	4		5	1	3	6	1
batman	126	5	6	4	4	1	7	7	5	5	1	3	5	8	4	5	8	6	4	4			7	9	5	4	5	4
korea	88	7	10	6	8	8	5	7	6	7	3	3	2		2		1	1	1	2		1		1	1	2	2	2
harv	119	3	5	4	4	1	4	6	7	5	4	9	7	5	4	6	9	3	2		2	2	4	2	11	4	4	2
ustutt	142	5	4	8	8	11	7	7	8	10	4	3	13	5	9	3	7	1		1	1	6	4	1	7	1	5	3
umann	185	12	11	6	12	4	10	9	10	9	13	4	11	11	9	7	3		8	4	2	8	3	2	5	4	4	4
umont	168	7	6	11	5	10	3	8	14	10	5	8	7	11	8	10		5	7	2	2	5	7	4	5	3	3	2
oce	216	20	10	11	3	8	4	17	17	8	23	9	11	10	10		13	4	6	3	4	6	2	4	2	2	2	7
ncar	225	9	16	15	8	10	14	13	14	16	8	16	20	7		4	5	6	5	4	2	6	5	3	5	4	2	8
connix	222	15	17	13	12	5	7	13	14	14	10	14	11		6	9	12	8	2	8	6	5	3	2	5	4	3	4
mit	200	14	10	13	11	8	6	9	5	10	13	10		8	6	13	7	10	11	6	2	5	5	6	4	4	2	2
sdsc	233	11	16	11	14	10	9	21	10	19	16		12	16	8	10	11	9	5	5	2	4	2	4		1	6	1
ukc	228	14	11	13	14	13	13	16	11	10		10	16	10	16	11	5	4	8	3	4	4	5	6	3	3	3	2
xor	302	24	20	22	19	22	16	27	21		14	14	18	11	8	11	5	9	4	4	5	6	6	3	5	4	1	3
wustl	308	25	23	27	20	16	16	20		24	17	18	14	13	14	9	11	2	4	6	5	5	3	7	4	1	3	1
bnl	374	34	31	30	35	26	32		32	25	16	14	9	6	13	14	9	8	3	5	7	4	3	3	1	4	3	7
ucl	431	41	44	40	45	43		33	24	20	14	16	15	7	11	11	7	6	3	7	9	3	7	5	6	8	2	4
lbli	506	56	51	72	59		41	31	28	26	13	27	7	13	16	6	10	7	6	3	10	2	4	5	4	2	4	3
ucol	436	47	53	45		45	44	26	28	14	9	16	10	13	23	6	4	8	3	6	14	5	5	2	2	7		1
nrao	454	57	48		52	42	44	28	32	24	12	14	8	15	12	9	4	7	6	2	8	5	5	4	3	3	7	3
usc	467	60		49	47	52	49	37	25	18	11	17	12	13	6	11	7	9	4	4	7	4	2	5	2	5	7	4
lbl	452		56	35	50	46	37	25	40	18	12	10	10	17	13	12	8	4	11	4	11	6	7	2	2	4	5	7
total	6459	501	473	456	456	396	385	382	380	319	252	251	241	220	214	192	166	145	124	115	109	109	106	103	100	93	92	79

Figure 5.3: Number of measurements made for each Internet path, $\mathcal{R}_1$ dataset

Destination	total	connix	lbl	ucla	sintef2	sintef1	austr2	inria	mid	rain	adv	ucl	near	wustl	umont	austr	ustutt	mit	sri	bsdi	bnl	umann	pubnix	sandia	sdsc	unij	ucol	harv	oce	ncar	nrao	lbli	ukc	panix
panix	682	13	15	23	9	14	17	32	18	23	32	13	16	41	16	45	27	12	29	16	13	20	15	24	14	20	28	34	19	27	23	11	23	
ukc	564	17	27	15	22	38	17	19	14	23	16	38	9	21	15	16	16	8	30	12	13	21	30	17	33	14	13	18	5	11	7	8		1
lbli	688	24	37	36	23	17	30	21	28	22	19	26	35	14	23	22	27	16	26	28	36	16	15	11	11	17	19	27	13	24	15		7	3
nrao	864	46	21	20	39	28	17	32	32	33	23	14	25	25	43	32	44	36	26	29	32	38	14	34	20	29	13	20	27	28		27	11	6
ncar	844	32	24	28	21	44	14	30	28	39	28	35	19	9	19	31	26	30	42	33	27	49	13	21	27	37	31	17	33		24	19	9	5
oce	1006	33	30	27	34	37	32	39	25	38	38	24	32	54	43	40	62	32	29	39	18	37	18	12	19	25	31	27		43	37	19	18	14
harv	931	31	34	30	34	38	30	44	32	26	44	41	20	31	20	26	23	23	32	37	37	13	40	41	35	36	23		26	15	13	30	11	15
ucol	921	27	34	31	45	29	34	32	31	13	22	24	46	26	32	32	35	25	40	44	26	53	31	21	16	34		17	33	31	12	23	14	8
unij	910	33	28	26	35	23	50	19	27	35	32	23	28	25	52	22	26	33	22	31	23	32	36	36	31		35	33	15	37	28	18	13	3
sdsc	1078	48	32	46	30	54	47	37	44	33	45	43	34	28	43	48	33	51	18	26	38	33	21	29		43	19	33	21	21	27	9	38	6
sandia	973	41	37	30	29	29	38	27	19	50	32	27	26	42	31	31	33	32	70	30	39	24	12		39	44	29	36	15	19	32	8	15	7
pubnix	1021	35	57	59	27	45	22	42	41	47	41	35	47	31	46	54	26	26	18	29	29	31		24	20	34	33	25	20	13	10	20	25	9
umann	1018	32	26	48	49	25	25	24	40	39	37	29	30	42	29	29	40	37	44	27	28		20	24	30	51	50	10	29	48	36	17	14	9
bnl	1031	37	28	41	36	49	86	28	31	23	31	39	40	44	32	33	34	30	38	18		27	28	32	33	22	24	46	18	27	29	32	12	3
bsdi	1043	45	42	29	60	43	54	26	28	28	43	34	31	51	24	37	16	23	30		18	37	37	39	28	32	27	30	39	35	27	29	10	11
sri	1043	43	35	27	34	19	31	47	38	34	37	32	32	42	44	21	39	21		26	39	40	15	64	17	28	40	36	28	36	21	26	31	20
mit	1030	39	41	38	48	43	55	32	82	42	40	41	24	27	28	34	18		21	21	29	37	26	32	52	32	25	22	15	25	33	15	7	6
ustutt	1110	36	36	60	33	33	39	34	37	39	37	44	39	51	41	23		18	41	21	33	41	25	34	28	40	37	22	61	26	41	29	23	8
austr	1065	42	41	36	51	46	46	30	35	48	37	14	32	36	41		17	36	23	37	35	36	51	31	43	23	28	32	38	31	28	8	9	24
umont	1077	26	24	72	49	34	20	25	34	38	36	47	39	43		55	42	30	32	36	28	29	31	27	33	59	31	19	34	19	41	26	13	5
wustl	1063	29	50	37	64	24	31	42	18	39	44	41	22		43	35	62	31	31	55	33	38	31	38	24	22	30	17	40	10	22	9	26	25
near	1228	81	46	68	44	48	42	54	48	51	29	57		40	54	34	43	43	34	33	36	26	48	35	34	25	40	17	30	27	28	21	9	3
ucl	1210	72	50	78	54	51	26	47	52	34	26		60	43	40	15	43	41	42	33	20	29	47	43	51	34	22	31	20	34	22	20	27	3
adv	1165	52	50	53	26	52	67	40	59	33		18	29	61	43	42	44	40	25	41	31	31	40	29	38	30	22	39	41	33	19	6	15	16
rain	1260	37	66	48	61	54	45	64	45		57	37	51	48	51	44	37	42	38	33	24	35	43	57	29	28	23	25	16	39	26	13	21	23
mid	1260	53	64	48	54	49	46	47		45	54	56	53	27	45	38	35	82	40	46	34	42	40	23	30	18	29	31	19	30	33	36	8	5
inria	1282	46	82	55	50	67	76		45	64	47	61	47	34	22	36	33	34	35	33	34	24	46	30	21	22	30	47	39	31	16	44	15	16
austr2	1168	51	49	22	45	42		66	53	44	67	21	44	31	18	41	56	54	37	46	75	26	29	30	37	30	23	38	28	12	14	15	13	11
sintef1	1336	87	61	32	43		61	65	51	41	67	37	46	34	42	43	35	46	16	53	56	29	48	22	53	35	26	27	43	44	23	15	44	11
sintef2	1277	57	59	45		31	42	43	44	69	29	71	42	68	39	41	39	48	47	40	38	46	22	29	39	48	31	28	35	21	19	27	27	13
ucla	1287	43	58		46	32	26	43	51	41	52	79	59	36	66	38	53	39	29	31	49	39	53	25	45	39	27	30	27	28	21	40	24	18
lbl	1309	51		47	57	61	49	81	54	61	28	46	46	44	43	51	28	41	47	38	39	32	59	51	32	13	30	41	27	25	14	47	21	5
connix	1365		52	43	36	87	64	43	53	36	52	68	88	34	26	43	39	39	60	37	44	33	47	41	36	25	44	30	33	33	53	27	11	8
total	35109	1339	1336	1298	1288	1286	1279	1255	1237	1231	1222	1215	1191	1183	1154	1132	1131	1099	1092	1059	1054	1044	1031	1006	998	989	913	905	887	883	794	694	564	320

Source

Figure 5.4: Number of measurements made for each Internet path, $\mathcal{R}_2$ dataset

Site	Best Guess
`wvnet-wtn9-c1.sura.net`	Charleston, WV
`128.167.205.2`	Charlottesville, VA
`reynolds-ctv1-c1.sura.net`	Charlottesville, VA
`uva-ctv-c3mb.sura.net`	Charlottesville, VA
`38.2.213.16`	New York, NY
`core.net218.psi.net`	New York, NY
`leaf.net218.psi.net`	New York, NY
`38.1.2.14`	Washington, D.C.
`core.net222.psi.net`	Washington, D.C.
`137.209.1.1`	College Park, MD
`192.80.6.2`	College Park, MD
`198.25.80.1`	College Park, MD
`199.54.78.1`	College Park, MD

Table IV: Uncertain router sites

Figures 5.3 and 5.4 summarize the number of `traceroute` measurements between each pair of sites for each of the experiments.

5.3 Geography

To understand the Internet topology traversed by the experiment, and how each router relates to others, we undertook to identify the geographic locations of the 751 routers (distinct IP addresses) involved in $\mathcal{R}_1$ and the 1,095 routers in $\mathcal{R}_2$. The identification involved several steps:

1. Routers with an Internet hostname in the same domain as one of the participating sites (e.g., `colorado.edu`) were assumed to be located at that site.

2. Routers with a single geographic location in their name (e.g., `dallas1.tx.alter.net`) were assumed to reside at that location.

3. For still-unidentified routers, we sent email to the NIC "whois" contacts [HSF85] for the router's domain, asking if they could identify the router's location or the naming scheme used for routers in that domain. The various contacts proved remarkably helpful, willing to go to considerable efforts to aid in locating the sites. We also benefited from various "whois" servers, especially the European server `whois.ripe.net` and its corresponding WAIS server, and topology maps.

4. If any still-unidentified routers only occurred as a hop between two identified sites at the same location, we assumed the router was sited at that location too. For example, if we observed a partial network path of $A \to B \to C$, with A and C both sited in San Diego, then we assumed that B was sited in San Diego too.

5. For the remainder, we made a "best guess," based on the locations of upstream and downstream routers. Table IV summarizes the sites for which we had to guess.

Thus, of the 1,531 routers traversed during the study, we were able to identify the location of all but 13.

After locating the routers, we reduced the topology traversed by the experiment to connections between cities, listed in Table V. Having developed a geographic database for the various routers, we then constructed maps showing the links traversed in the study.[3] Figure 5.5 shows these links from a North American perspective, where sites in Hawaii, Korea, and Australia are shown west of California, and sites in Europe and Israel are shown in the Atlantic. Figure 5.6 show the links from a European perspective; here, the only links extending outside of Europe were those to sites in the U.S., which is represented as a single site west of France.

[3] Doing so first required analyzing the `traceroutes` for routing pathologies (§ 6), because "fluttering" and mid-stream routing changes can easily introduce spurious links.

State or Country	City
California	Anaheim, Berkeley, Bloomington, Hayward, Livermore, Los Angeles, NASA-AMES (Moffett Field), Oakland, Palo Alto, Pasadena, Sacramento, San Diego, San Francisco, San Jose, Santa Clara, Stockton
Colorado	Boulder, Colorado Springs, Denver, East Boulder
Connecticut	Hartford, Middlefield
Florida	Miami
Georgia	Atlanta
Hawaii	Honolulu
Illinois	Batavia, Chicago, Willow Springs
Maryland	College Park
Massachusetts	Boston, Cambridge, Waltham
Michigan	Detroit
Missouri	Kansas City, St. Louis
Nebraska	Lincoln
New Jersey	Pennsauken, Princeton, West Orange
New Mexico	Albuquerque, Los Alamos
New York	Albany, Armonk, Brookhaven, Buffalo, Deer Park, Ithaca, New York, Syracuse
North Carolina	Greensboro, Raleigh
Ohio	Cleveland, North Royalton
Oregon	Portland
South Carolina	Greenville
Texas	Austin, Dallas, Fort Worth, Houston
Virginia	Charlottesville, Fairfax, Falls Church, Newport News, Norfolk, Vienna
Washington, D.C.	
Washington	Kent, Seattle
West Virginia	Charleston
Australia	Adelaide, Canberra, Melbourne, Newcastle, Sydney
Austria	Vienna
Belgium	Brussels
Canada	Vancouver, Montreal, Toronto
England	Cambridge, Canterbury, London, Manchester
Finland	Helsinki
France	Lyon, Marseilles, Montpellier, Nice, Paris, Poitiers, Sophia, Toulouse
Germany	Aachen, Duesseldorf, Heidelberg, Karlsruhe, Mannheim, Munich, Stuttgart
Italy	Milan
Israel	Jerusalem, Rehovot
Korea	Pohang, Seoul
Netherlands	Amersfoort, Amsterdam, Den Bosch, Eindhoven, Nijmegen, Venlo, Utrecht
Norway	Oslo, Trondheim
Spain	Madrid
Sweden	Stockholm
Switzerland	Geneva

Table V: Router cities

Figure 5.5: Links traversed during $\mathcal{R}_1$ and $\mathcal{R}_2$, North American perspective

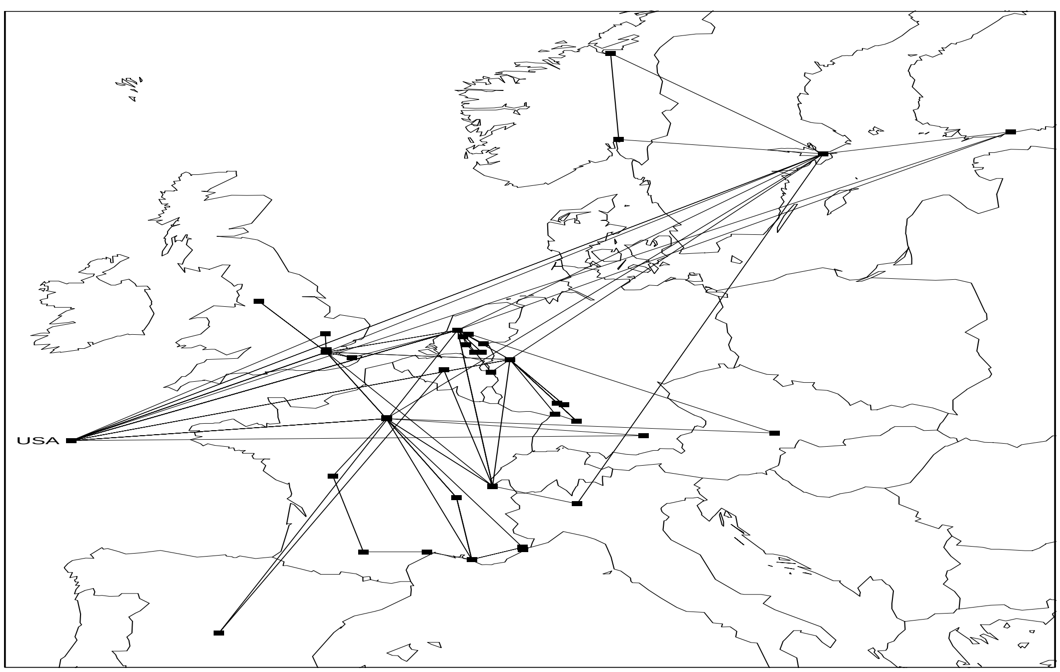

Figure 5.6: Links traversed during $\mathcal{R}_1$ and $\mathcal{R}_2$, European perspective