

Contents

Acknowledgments	xxi
I Background	1
1 Introduction	3
1.1 Why Measure the Internet?	5
1.2 How to Read this Book	6
1.3 Resources for More Information	10
2 Internet Architecture	13
2.1 The Internet's Architecture	13
2.1.1 The History of the Internet	14
2.1.2 The Organization of the Internet	17
2.1.3 Design Principles of the Internet	22
2.2 Details of Internet Operation	25
2.2.1 Endsystems, Links, and Routers	25
2.2.2 Autonomous Systems	26
2.2.3 Routing	26
2.3 Protocols	31
2.3.1 IP	32
2.3.2 TCP	33
2.3.3 UDP	34
2.3.4 Routing Protocols	35
2.3.5 ICMP	35
2.3.6 SNMP	36
2.3.7 IP Multicast	37
2.3.8 DNS	37
2.3.9 HTTP	39

2.3.10	P2P	40
2.4	Applications	42
3	Analytic Background	45
3.1	Linear Algebra	45
3.2	Probability	48
3.2.1	Background	49
3.2.2	Special Issues in the Internet	55
3.3	Statistics	58
3.3.1	Background	58
3.3.2	Special Issues in the Internet	61
3.4	Graphs	64
3.4.1	Background	64
3.4.2	Special Issues in the Internet	66
3.5	Metrics	70
3.6	Measurement and Modeling	73
3.6.1	Models in General	73
3.6.2	The Use of Probability Models	76
4	Practical Issues in Internet Measurement	79
4.1	Where can Measurements be Made?	80
4.1.1	Local Area Network	82
4.1.2	Inside a Backbone	82
4.1.3	Entry Points into a Network	84
4.1.4	Mirror Sites/Network Exchange Points	86
4.1.5	Wide Area Network	87
4.2	Role of Time	89
4.2.1	Background	90
4.2.2	Sources of Time Information	91
4.2.3	Synchronized Time	93
4.3	Role of Internet Directories and Databases	94
4.3.1	Internet Address and Routing Registries	95
4.3.2	Domain Name System	97
4.3.3	Measurement-related Issues in Dealing with Databases	98
4.4	Measurements Across Various Protocol Layers	99
4.4.1	Issues in Capturing Data	99
4.4.2	Changes to Infrastructure/Instrumentation	102
4.4.3	Local vs. Remote vs. Distributed Data Gathering	103
4.4.4	Measurement on Overlays	104

II	In Depth	105
5	Infrastructure	107
5.1	Properties	107
5.1.1	Physical Device Properties	107
5.1.2	Topology Properties	111
5.1.3	Interaction of Traffic and Network	112
5.2	Challenges	115
5.2.1	Core Simplicity	115
5.2.2	Hidden Layers	116
5.2.3	Hidden Pieces	116
5.2.4	Administrative Barriers	117
5.3	Tools	117
5.3.1	Active Measurement	118
5.3.2	Passive Measurement	124
5.3.3	Fused Measurements	127
5.3.4	Bandwidth Measurement	127
5.3.5	Latency Measurement and Estimation	136
5.3.6	Geolocation	142
5.3.7	Inference	147
5.3.8	Other Tools	152
5.4	State of the Art	152
5.4.1	Equipment Properties	153
5.4.2	Topology Properties	154
5.4.3	Interaction of Traffic and Network	165
6	Traffic	171
6.1	Properties	172
6.1.1	The Basics: Packets and Bytes	172
6.1.2	Higher-level Structure	173
6.1.3	Flows	175
6.1.4	Semantically Distinct Traffic Types	176
6.2	Challenges	176
6.2.1	Practical Issues	177
6.2.2	Statistical Difficulties	179
6.3	Tools	188
6.3.1	Packet Capture	188
6.3.2	Data Management	191
6.3.3	Data Reduction	192
6.3.4	Inference	212

6.4	State of the Art	215
6.4.1	Packets and Bytes	216
6.4.2	Higher-level Structure	234
6.4.3	Flows	236
6.4.4	Control Traffic	238
6.4.5	Wireless	239
7	Applications	241
7.1	Application Mix	242
7.2	DNS	244
7.2.1	DNS Measurement Properties	245
7.2.2	DNS Measurement Challenges	248
7.2.3	DNS Measurement Tools	251
7.2.4	Use of DNS in Other Applications	256
7.2.5	State of the Art	258
7.3	Web	269
7.3.1	Web Measurement Properties	270
7.3.2	Web Measurement Challenges	273
7.3.3	Web Measurement Tools	278
7.3.4	State of the Art	286
7.4	P2P	309
7.4.1	P2P Measurement Properties	310
7.4.2	P2P Measurement Challenges	314
7.4.3	P2P Measurement Tools	317
7.4.4	State of the Art	321
7.5	Online Games	331
7.5.1	Games and Measurement Properties	332
7.5.2	Networked Games Measurement Challenges	337
7.5.3	State of the Art	340
7.6	Other Applications	346
7.6.1	Streaming Multimedia	346
III	In Perspective	353
8	Anonymization	355
8.1	Definitions	356
8.2	General Motivation for Anonymizing Data	357
8.3	Obstacles and Risks in Sharing Data	358
8.4	What Should be Anonymized: Data Categorization	360

8.5	How Data is Anonymized: Process and Techniques	365
8.5.1	Anonymization Process	365
8.5.2	Anonymization Techniques	367
8.6	Anonymization Examples at Different Layers	369
8.6.1	Configuration Data	369
8.6.2	Router-level Data	370
8.6.3	Packet-level Traces	370
8.6.4	Application-level Data	373
8.7	Attacks Against Anonymized Data	374
8.8	Anonymizing Data: Metrics for Success	376
8.9	Alternatives to Anonymization	377
9	Security	379
9.1	Role of Internet Measurement in Security	380
9.2	Intranet Measurements in Aid of Security	382
9.3	Gateway Measurements in Aid of Security	384
9.4	Inter-domain Measurements Impact on Security	386
9.5	Wide-area Measurements in Aid of Security	387
9.6	Application-level Measurements of Attacks	394
10	Case Studies	395
10.1	Low-level Monitoring Tools	395
10.2	Individual Toolsets for Network Measurement	397
10.2.1	Windmill	398
10.2.2	Click	399
10.2.3	dss	400
10.2.4	Gigascope	403
10.3	Large-scale Measurement Projects	404
10.3.1	RIPE	405
10.3.2	High-energy Physics	407
10.3.3	CAIDA	410
10.3.4	PlanetLab	414
11	Conclusions and Prospects	419
11.1	Trends in Internet Measurement	419
11.2	Difficulties	424
11.3	Future Work	426
11.3.1	Research Challenges	426
11.3.2	Emerging Questions	428

Bibliography	431
Index	473

List of Tables

2.1	Common link technologies and speeds	27
3.1	Notation	46
3.2	Distributions commonly encountered in Internet modeling	52
3.3	Model properties and data properties	77
5.1	Variation in People/Interface Density Across Regions	165
6.1	Fields in the <code>ifEntry</code> of the MIB-II interfaces group	194
7.1	DNS properties of interest to measure.	245
7.2	DNS measurement tools.	252
7.3	Using DNS for other applications.	257
7.4	Web properties of interest to measure.	271
7.5	Notable P2P protocols.	310
7.6	P2P properties of interest to measure.	315
7.7	Networked games properties of interest to measure.	333
8.1	Need for anonymization: Attack techniques and targets.	359
8.2	Anonymizable data categorization per-protocol/application.	362
8.3	Anonymization process.	366
8.4	Anonymization techniques: transformation.	367
8.5	IP header fields to be anonymized.	371
10.1	Example tools developed or supported at CAIDA.	411

List of Figures

1.1	Organization of the book	8
2.1	A packet-switched network	18
2.2	Protocol layers in the Internet	19
2.3	Nested headers	20
2.4	The IP hourglass	23
2.5	Hierarchical organization and routing	29
3.1	Comparison of short- and long-tailed distributions	57
3.2	Example histogram and empirical CDF	60
3.3	Distribution of categorical data: frequency vs. rank	61
3.4	Example of Zipf's law	63
3.5	Example graphs: undirected and directed	64
3.6	Example graphs: regular and random	69
4.1	Measurement locations at an ISP.	81
4.2	IPv4 address distribution.	96
5.1	Typical router organization	108
5.2	Internals of the forwarding engine	109
5.3	Traceroute in operation	119
5.4	Traceroute behavior with unstable paths	120
5.5	Partial AS topology from BGP	125
5.6	Packet-pair spacing set at narrow link	130
5.7	The probe gap method	131
5.8	Multilateration with geographic distance constraints	145
5.9	The tomography setting	148
5.10	Highly variable degree distribution in the AS graph	154
5.11	Map of the ARPANET as of December 30, 1972	158
5.12	Map of the Abilene network as of 2005	159
5.13	Example synthetic graph meeting technological constraints for routers	160

5.14	Number of unique AS numbers advertised within the BGP system	161
5.15	Number of hosts registered in DNS	162
5.16	Semi-log plot of hosts registered in DNS	163
6.1	Three views of a traffic trace	173
6.2	Three levels of structure in traffic	174
6.3	Running means and standard deviations of highly variable data	180
6.4	Total size measure as a function of smallest observations	183
6.5	Traffic at two timescales: unstable vs. stable	184
6.6	Effects of autocorrelation on confidence intervals	187
6.7	Example query in Gigascope	192
6.8	An example of a Bloom filter	200
6.9	Illustration of PCA on a correlated, 2-D dataset	205
6.10	Scree plot for origin–destination flows	207
6.11	Fitted distribution to traffic data from Figure 6.5(b)	208
6.12	Visual checks for goodness-of-fit	210
6.13	Autocorrelation function for traffic data from Figure 6.5(b)	212
6.14	Flows in a network	214
6.15	Overview of traffic analysis	216
6.16	One week of traffic on two links of the Abilene network	217
6.17	Autocorrelation function of byte counts $\{B_n\}$	220
6.18	Traffic scaling properties	223
6.19	Packet trains and autocorrelation	225
6.20	LLCD of file sizes on Unix systems	226
6.21	Logscale diagram for traffic data from Figures 6.17 and 6.18(a)	228
7.1	DNS entities and transactions.	249
7.2	Converting <i>netflow</i> data into a graph.	260
7.3	King tool operation.	264
7.4	DNS-Enhanced Web architecture.	266
7.5	Clients with varying connectivity accessing a Web server complex.	275
7.6	Multiple servers and layers involved in fetching a resource.	277
7.7	Time-series of counts of the requests in a server log: (a) the entire server log; (b) client cluster with a proxy; (c) client cluster with a spider. From [KW00].	285
7.8	Distribution of clients and clusters over time. From [JKR02].	304
7.9	The number of (a) requests per second, (b) number of clients vs. clusters in a second, in an attack victim trace. From [JKR02].	305
7.10	The number of links and unique links	308
7.11	The number of self-references	309

7.12 BitTorrent seeds, leechers, and torrent file.	328
7.13 BitTorrent client downloading from a peer set.	329
9.1 Mohonk architecture.	389
10.1 <i>dss</i> component architecture.	401

Acknowledgments

We would like to thank many people for their time, interest, and willingness to help us with the book. The reviewers are not responsible for any errors in the book. We thank them for answering questions, reading draft versions of the chapters, and giving valuable feedback.

We thank Carey Williamson for his review of the entire book and numerous thoughtful comments. Our thanks to Greg Minshall for his complete read and for providing valuable feedback on a number of topics.

We also thank Paul Barford and Dina Papagiannaki for their reviews of several chapters and useful suggestions. We also thank the various reviewers who reviewed on behalf of the publisher including Constantine Dovrolis, Chen-Nee Chuah, Russell Clark, Yan Chen, and Anat Bremler Barr.

The comments of David Poole (Chapter 3), Supratik Bhattacharyya (Chapter 4), Constantine Dovrolis (Chapters 5, 6, and 10), Walter Willinger and Matthew Roughan (Chapter 6), Anees Shaikh, Michael Rabinovich, and Craig Wills (Chapter 7), Duane Wessels (Chapter 7, on DNS), Yatin Chawathe (Chapter 7, on P2P), Jacobus van der Merwe (Chapter 7, on multimedia), Mark Claypool (Chapter 7, on networked games), Greg Minshall, Lorrie Cranor, Martin Arlitt, and Jeffrey Mogul (Chapter 8), Mark Allman and Henk Uijterwaal (Chapter 10) were very helpful. They took time off their busy schedules to give detailed feedback electronically and in many cases face to face. We thank them for the same.

We thank Jay Borkenhagen for his insights on a variety of topics related to evolution of ISPs and their internal engineering and Kavé Salamatian for discussions concerning the nature of Internet modeling. Wu-Chang Feng's detailed comments on the networked games section were very helpful. Duane Wessels provided valuable help on DNS and answered several questions.

We also thank Roy Arends, Martin Arlitt, Grenville Armitage, Steven Bellovin, Andre Broido, Nevil Brownlee, Benoit Claise, Edith Cohen, Michalis Faloutsos, Chris Frazier, Emden Gansner, Tristan Henderson, Sugih Jamin, Leonard Kleinrock, Eric Kolaczyk, George Kollis, Anukool Lakhina, Carsten Lund, David Meyer, Jörg Micheel, Mikhail Mikhailov, Michael Mitzenmacher, Mark Nottingham,

Andrew Odlyzko, Jitendra Padhye, Dave Plonka, Jürgen Quittek, Sylvia Ratnasamy, Matthew Roughan, Henning Schulzrinne, Aman Shaikh, Oliver Spatscheck, Paul Vixie, Jim Xu, and Artur Ziviani for answering various questions.

We would like to thank Jonathan Shipley at Wiley for shepherding our book through the editorial process. We also thank Vivian Ward, Deborah Egleton, Sam Crowe, Claire Jardine, David Barnard, and Sarah Lewis for their help in processing the book. Our thanks also to Gaynor Redvers-Mutton who assisted us early on.

Mark wrote portions of this book at Boston University with support from Sprint, Intel, and the National Science Foundation (NSF), and portions while he was on sabbatical at the Laboratoire d'Informatique de Paris 6 (LIP6) with support from Centre National de la Recherche Scientifique (CNRS). Both institutions are exciting intellectual environments and wonderful places to work; I thank my colleagues on both sides of the Atlantic for their support and encouragement. Most importantly, I have depended on my family and friends at many times and in many ways as I wrote this book and I want to thank them (you know who you are!) from the bottom of my heart.

Balachander would like to thank AT&T Labs–Research management, especially David Belanger, for his encouragement. Thanks to tnn for keeping me supplied with my drugs of choice: r adams, j farrar, and rare bruce juice. Thanks to Karthik L for the carnatica cornucopia. It was a pleasure to partner with Wendy on the cover design. Thanks to mjs for being an able attorney and a better friend, albeit one suffering from a pollyannish optimism of ever surpassing me in the ranks of bumhood. Balachander wrote this book on an IBM Thinkpad T-41P running Linux, in several places around the world; the hosts in various places deserve thanks including Chris and Mary Malley, Rob Mason; special thanks to Cristina Ruggieri for the extended use of her mother's house south of Il colle Aventino.