

Table of Contents

1. Introduction	1
1.1 Why We Wrote This Book	1
1.2 Structure and Contents	3
1.3 Teaching Materials and Book Website	5
1.4 Acknowledgements	5

Part I. Peer-to-Peer: Notion, Areas, History and Future

2. What Is This “Peer-to-Peer” About?	9
2.1 Definitions	10
2.1.1 Shift of Paradigm in Internet Communication	12
2.2 Research Challenges in Peer-to-Peer Systems & Applications .	12
2.2.1 Unstructured Peer-to-Peer Systems	15
2.2.2 Structured Peer-to-Peer Systems	15
2.3 Conclusion	16
3. Past and Future	17
3.1 Status Quo: Networks (Over)Filled with Peer-to-Peer Traffic .	17
3.2 How It All Began: From Arpanet to Peer-to-Peer	18
3.3 The Napster-Story	19
3.4 Gnutella and Its Relatives: Fully Decentralized Architectures .	20
3.5 Driving Forces Behind Peer-to-Peer	22
4. Application Areas	25
4.1 Information	25
4.2 Files	27
4.3 Bandwidth	29
4.4 Storage Space	30
4.5 Processor Cycles	31

Part II. Unstructured Peer-to-Peer Systems

5. First and Second Generation of Peer-to-Peer Systems	35
5.1 General Characteristics of Early Peer-to-Peer Systems	35
5.2 Centralized Peer-to-Peer Networks	37
5.2.1 Basic Characteristics	37
5.2.2 Signaling Characteristics	38
5.2.3 Discussion	41
5.3 Pure Peer-to-Peer-Networks	42
5.3.1 Basic Characteristics	42
5.3.2 Signaling Characteristics	44
5.3.3 Discussion	46
5.4 Hybrid Peer-to-Peer Networks	49
5.4.1 Basic Characteristics	49
5.4.2 Signaling Characteristics	52
5.4.3 Discussion	54
6. Random Graphs, Small-Worlds and Scale-Free Networks .	57
6.1 Introduction	57
6.2 Definitions	59
6.3 The Riddle – Analysis of Real Networks	60
6.4 Families and Models	61
6.4.1 Random Graphs	61
6.4.2 Small-Worlds – The Riddle’s First Solution	64
6.4.3 Scale-Free Networks: How the Rich Get Richer	67
6.5 Applications to Peer-to-Peer Systems	70
6.5.1 Navigating in Small-Worlds	70
6.5.2 Small-World Overlay Networks in P2P Systems	72
6.5.3 Scale-Free Overlay Networks in P2P Systems	75
6.6 Summary	76

Part III. Structured Peer-to-Peer Systems

7. Distributed Hash Tables	79
7.1 Distributed Management and Retrieval of Data	80
7.1.1 Comparison of Strategies for Data Retrieval	81
7.1.2 Central Server	81
7.1.3 Flooding Search	82
7.1.4 Distributed Indexing – Distributed Hash Tables	84
7.1.5 Comparison of Lookup Concepts	85
7.2 Fundamentals of Distributed Hash Tables	86
7.2.1 Distributed Management of Data	86
7.2.2 Addressing in Distributed Hash Tables	86

7.2.3	Routing	88
7.2.4	Data Storage	89
7.3	DHT Mechanisms	89
7.3.1	Overview	90
7.3.2	Node Arrival	90
7.3.3	Node Failure	90
7.3.4	Node Departure	91
7.4	DHT Interfaces	91
7.4.1	Routing Interface	92
7.4.2	Storage Interface	92
7.4.3	Client Interface	92
7.5	Conclusions	93
8.	Selected DHT Algorithms	95
8.1	Chord	95
8.1.1	Identifier Space	95
8.1.2	Routing	96
8.1.3	Self-Organization	97
8.2	Pastry	99
8.2.1	Identifier Space	100
8.2.2	Routing Information	100
8.2.3	Routing Procedure	102
8.2.4	Self-Organization	102
8.2.5	Routing Performance	105
8.3	Content Addressable Network <i>CAN</i>	106
8.3.1	Identifier Space	107
8.3.2	Routing Information	108
8.3.3	Routing Procedure	109
8.3.4	Self-Organization	109
8.3.5	Routing Performance	111
8.4	Symphony	112
8.5	Viceroy	113
8.6	Kademlia	114
8.7	Summary	116
9.	Reliability and Load Balancing in Distributed Hash Tables	119
9.1	Storage Load Balancing of Data in Distributed Hash Tables	119
9.1.1	Definitions	121
9.1.2	A Statistical Analysis	121
9.1.3	Algorithms for Load Balancing in DHTs	124
9.1.4	Comparison of Load-Balancing Approaches	129

9.2	Reliability of Data in Distributed Hash Tables	131
9.2.1	Redundancy	132
9.2.2	Replication	132
9.3	Summary	135
10.	P-Grid: Dynamics of Self-Organizing Processes in Structured Peer-to-Peer Systems	137
10.1	The Concept of Self-Organization	137
10.2	Example of Self-Organization in Unstructured P2P Systems ..	138
10.3	Self-Organization in Structured Peer-to-Peer Systems	140
10.3.1	The Structure of P-Grid Overlay Networks	141
10.3.2	Dynamics of P-Grid Overlay Networks	143
10.3.3	Bootstrapping a P-Grid Overlay Network	144
10.3.4	Routing Table Maintenance	146
10.3.5	Analysis of the Maintenance Mechanism	150
10.4	Summary	151
<hr/>		
Part IV. Peer-to-Peer-Based Applications		
<hr/>		
11.	Application-Layer Multicast	157
11.1	Why Multicast on Application Layer	157
11.2	Design Aspects and Taxonomy	158
11.3	Unstructured Overlays	159
11.3.1	Centralized Systems	159
11.3.2	Fully Distributed Systems	161
11.4	Structured Overlays	163
11.4.1	Flooding-Based Replication	164
11.4.2	Tree-Based Replication	165
11.4.3	Performance/Cost Evaluation	168
11.5	Hot Topics	169
11.6	Summary	170
12.	ePOST	171
12.1	Scoped Overlays	172
12.1.1	Design	173
12.1.2	Ring Structure	173
12.1.3	Gateway Nodes	175
12.1.4	Routing	175
12.1.5	Global Lookup	176
12.2	POST Design	176
12.2.1	Data Types	177
12.2.2	User Accounts	178
12.2.3	Single-Copy Store	179
12.2.4	Event Notification	179

12.2.5	Metadata	180
12.2.6	Garbage Collection	181
12.2.7	POST Security	182
12.3	ePOST Design	184
12.3.1	Email Storage	184
12.3.2	Email Delivery	184
12.3.3	Email Folders	185
12.3.4	Incremental Deployment	186
12.3.5	Management	186
12.4	Correlated Failures	187
12.4.1	Failure Models	188
12.4.2	Glacier	189
12.4.3	Maintenance in Glacier	190
12.4.4	Recovery After Failures	191
12.4.5	Object Aggregation	191
12.5	Preliminary Experience	192
13.	Distributed Computing – GRID Computing	193
13.1	Introduction	193
13.2	The GRID Architecture	194
13.3	The Globus Project	196
13.4	Defining the GRID: The Global GRID Forum Initiative	198
13.4.1	The Open GRID Services Architecture (OGSA)	198
13.4.2	GRID Services: Building Blocks for the GRID	200
13.4.3	Stateful Web Services: OGSF & WS-Resource FW	201
13.5	GRID and Peer-to-Peer Computing	202
13.5.1	Comparing GRID and Peer-to-Peer	203
13.5.2	GRID and Peer-to-Peer: Converging Concepts?	204
13.6	Summary	205
14.	Web Services and Peer-to-Peer	207
14.1	Introduction	207
14.2	Architecture and Important Standards	209
14.2.1	XML and XML Schema	211
14.2.2	WSDL	212
14.2.3	SOAP	216
14.2.4	HTTP	217
14.2.5	UDDI	217
14.2.6	WS-*	217
14.3	Service Orchestration	219
14.4	Comparison of Peer-to-Peer and Web Services	219
14.4.1	What Can Peer-to-Peer Learn from Web Services?	220
14.4.2	What Can Web Services Learn from Peer-to-Peer?	221
14.4.3	Side-Effects when Joining Web Services and P2P	222
14.5	Resulting Architectures	223

Part V. Self-Organization

15. Characterization of Self-Organization	227
15.1 Introduction	227
15.2 Basic Definitions	228
15.2.1 System	228
15.2.2 Complexity	229
15.2.3 Feedback	230
15.2.4 Emergence	230
15.2.5 Complex System	231
15.2.6 Criticality	232
15.2.7 Hierarchy & Heterarchy	233
15.2.8 Stigmergy	235
15.2.9 Perturbation	235
15.3 Characteristics of Self-Organization	235
15.3.1 Self-Determined Boundaries	235
15.3.2 Operational Closure & Energetic Openness	236
15.3.3 Independence of Identity and Structure	236
15.3.4 Maintenance	237
15.3.5 Feedback & Heterarchy	237
15.3.6 Feedback	237
15.3.7 Criticality	238
15.3.8 Emergence	238
15.3.9 Self-Determined Reaction to Perturbations	239
15.3.10 Reduction of Complexity	239
15.4 Applications in Computer Science	239
15.4.1 Small-World and Scale-Free Networks	240
15.4.2 Swarming	242
15.4.3 Cellular Automata	244
15.5 Conclusions	245
16. Self-Organization in Peer-to-Peer Systems	247
16.1 Introduction	247
16.2 Evaluation of Peer-to-Peer Systems	248
16.2.1 Criteria	248
16.2.2 Unstructured Peer-to-Peer Networks	250
16.2.3 Structured Peer-to-Peer Systems	255
16.2.4 Summary of Peer-to-Peer Evaluations	259
16.3 Towards More Self-Organization in Overlays	260
16.3.1 Active Virtual Peers	260
16.3.2 Objectives & Requirements for Control of Overlays ...	261
16.3.3 An Implementation of the AVP Concept	262
16.3.4 Related Work	265
16.4 Conclusions	265

Part VI. Search and Retrieval

17. Peer-to-Peer Search and Scalability	269
17.1 Peer-to-Peer Search and Lookup in Overlay Networks	269
17.1.1 Problem Statement and Chapter Overview	270
17.1.2 Search and Lookup – Functional Options	270
17.1.3 Design Space	272
17.1.4 Overlay Topology Requirements	274
17.1.5 Overlay Topology Parameters	275
17.2 Scalability in Peer-to-Peer Systems	276
17.2.1 Definition of Peer-to-Peer Scalability	277
17.2.2 Efficiency and Scale	278
17.2.3 Scalability Metric and Notation	280
17.3 A Scheme for Lookup and Search Overlay Scalability	281
17.3.1 Overhead for Lookup and Search	282
17.3.2 Dimensions of Lookup & Search Overhead	283
17.3.3 The Assessment Scheme	284
17.4 Scalable Search with SHARK	285
17.5 Summary and Conclusions	288
18. Algorithmic Aspects of Overlay Networks	289
18.1 Background and Motivation	289
18.2 Model Definition	293
18.3 Gathering Information Along a Path	295
18.3.1 Basic Algorithms	295
18.3.2 collect-rec	297
18.4 weighted collect-rec Algorithm	300
18.4.1 Algorithm Description	302
18.4.2 Detailed Algorithm Description	302
18.4.3 Analysis of weighted collect-rec Algorithm	305
18.5 Gathering Information from a Tree	307
18.5.1 Detailed Algorithm Description	309
18.5.2 Analysis of weighted collect on trees Algorithm	311
18.6 Gathering Information from General Graphs	315
18.7 Global Functions	315
18.8 Performance Evaluation	316
18.8.1 weighted collect-rec Algorithm Performance	316
18.8.2 Performance of weighted collect on trees Algorithm	319
19. Schema-Based Peer-to-Peer Systems	323
19.1 Introduction	323
19.2 Design Dimensions of Schema-Based Peer-to-Peer Systems	325
19.2.1 Data Model and Query Language	325

19.2.2	Data Placement	326
19.2.3	Topology and Routing	327
19.3	Case Study: A Peer-to-Peer Network for the Semantic Web . .	329
19.3.1	Semantic Web Data Model and Query Language	329
19.3.2	Schema-Based Routing Indices	331
19.4	Advanced Topics	333
19.4.1	Schema Mapping	333
19.4.2	Distributed Query Plans	334
19.4.3	Top- k Query Processing	334
19.5	Conclusion	336
20.	Supporting Information Retrieval in Peer-to-Peer	
	Systems	337
20.1	Content Searching in Peer-to-Peer Applications	337
20.1.1	Exchanging Media Files by Meta-Data Searches	338
20.1.2	Problems in Peer-to-Peer Information Retrieval	338
20.1.3	Related Work in Distributed Information Retrieval . . .	341
20.2	Indexstructures for Query Routing	343
20.2.1	Distributed Hash Tables for Information Retrieval . . .	344
20.2.2	Routing Indexes for Information Retrieval	345
20.2.3	Locality-Based Routing Indexes	347
20.3	Supporting Effective Information Retrieval	348
20.3.1	Providing Collection-Wide Information	348
20.3.2	Estimating the Document Overlap	349
20.3.3	Prestructuring Collections with Taxonomies	350
20.4	Summary and Conclusion	351
21.	Hybrid Peer-to-Peer Systems	353
21.1	Introduction	353
21.2	Overlay Network Design Dimensions	354
21.3	Hybrid Architectures	356
21.3.1	JXTA	357
21.3.2	Brocade	358
21.3.3	SHARK	360
21.3.4	Omicron	360
21.4	Hybrid Routing	363
21.4.1	OceanStore	363
21.4.2	Hybrid PIER	364
21.5	Comparison with Non-hybrid Systems	364
21.6	Summary and Conclusion	365

Part VII. Peer-to-Peer Traffic and Performance Evaluation

22. ISP Platforms Under a Heavy Peer-to-Peer Workload	369
22.1 Introduction	369
22.2 Peer-to-Peer Traffic Characteristics	370
22.2.1 Traffic Mix on IP Platforms	370
22.2.2 Daily Traffic Profile	370
22.2.3 Traffic Growth and Prognosis	373
22.2.4 Asymmetrical Versus Symmetrical Access Lines	373
22.3 Cross Layer Aspects	374
22.3.1 Routing on Application and IP Layer	374
22.3.2 Network and Transport Layer Analysis	375
22.3.3 Application Layer Pattern	375
22.3.4 Distribution of Sources for eDonkey File-Sharing	376
22.3.5 Caches for Peer-to-Peer Data	378
22.4 Implications for QoS in Multi-service IP Networks	379
22.5 Conclusion	380
23. Traffic Characteristics and Performance Evaluation of Peer-to-Peer Systems	383
23.1 Introduction	383
23.2 A Concept for Peer-to-Peer Performance	384
23.3 Traffic Characteristics of Peer-to-Peer-Systems	386
23.3.1 Gnutella	386
23.3.2 eDonkey	387
23.4 Evaluation of a Peer-to-Peer Resource Mediation Mechanism	391
23.5 Evaluation of a Peer-to-Peer Resource Access Mechanism	394
23.6 Conclusion	396

Part VIII. Peer-to-Peer in Mobile and Ubiquitous Environments

24. Peer-to-Peer in Mobile Environments	401
24.1 Why Is P2P Interesting for Mobile Users and Services	401
24.1.1 Scenario 1: Taxi Locator	402
24.1.2 Scenario 2: University Campus	403
24.2 Introduction to Mobile Communication Systems	403
24.3 Challenges for Peer-to-Peer Techniques in Mobile Networks	405
24.3.1 Peer-to-Peer Systems in Mobile Ad-Hoc Networks	406
24.4 Solutions for Peer-to-Peer in Mobile and Wireless Networks	407
24.4.1 Solutions with Unstructured Peer-to-Peer Networks	408
24.4.2 Solutions Based on Structured Peer-to-Peer Networks	412
24.5 Summary	416

25. Spontaneous Collaboration in Mobile Peer-to-Peer Networks	419
25.1 Introduction and Motivation	419
25.1.1 Mobile Peer-to-Peer Networks and MANETS	420
25.1.2 One-Hop Peer-to-Peer Design Space	422
25.1.3 Chapter Overview	423
25.2 Application Domains and Examples	423
25.2.1 Shark	423
25.2.2 MobiTip	424
25.2.3 SpotMe	424
25.2.4 Socialight	425
25.2.5 AdPASS	425
25.3 Building Blocks for Mobile Peer-to-Peer Networks	426
25.4 The iClouds Project	429
25.4.1 Multi-hop Information Dissemination	429
25.4.2 Data Structures and Communication Semantics	430
25.4.3 Architecture	432
25.5 Conclusion	433
26. Epidemic Data Dissemination for Mobile Peer-to-Peer Lookup Services	435
26.1 Motivation and Background	435
26.2 Passive Distributed Indexing	436
26.2.1 Overview	436
26.2.2 Basic Concept	437
26.2.3 Selective Forwarding for Extending Radio Coverage	439
26.3 Consistency Issues	440
26.3.1 Dealing with Weak Connectivity and Node Failures	440
26.3.2 Dealing with Data Modification at the Origin Node	441
26.4 Performance Studies	443
26.4.1 Simulation Environment	443
26.4.2 Sensitivity to System Characteristics	447
26.4.3 Sensitivity to Application Characteristics	450
26.4.4 Impact of Consistency Mechanisms	452
26.5 Summary	454
27. Peer-to-Peer and Ubiquitous Computing	457
27.1 Introduction to Ubiquitous Computing	457
27.2 Characteristics of Ubiquitous Computing Applications	458
27.2.1 Information	459
27.2.2 Network	459
27.2.3 Collaboration	460
27.2.4 Sharing Resources	460
27.2.5 Context Information	460
27.3 Communications in Ubiquitous Computing Architectures	461

27.4	Ubiquitous Computing Middleware	461
27.4.1	Support for Heterogeneous Devices	462
27.4.2	Resource Constraints	462
27.4.3	Mobility Support	462
27.4.4	Networking Support	463
27.4.5	Performance Issues	463
27.5	Peer-to-Peer and Ubiquitous Computing	465
27.6	Research Challenges in Ubiquitous Peer-to-Peer Computing	466
27.6.1	Heterogeneous Devices	467
27.6.2	Efficient Algorithms	467
27.6.3	Security and Privacy	467
27.6.4	Scalable Architectures	468
27.6.5	Next Generation Peer-to-Peer Middleware	468
27.7	Summary	468

Part IX. Business Applications and Markets

28.	Business Applications and Revenue Models	473
28.1	Introduction	473
28.2	Definitions	474
28.2.1	Peer-to-Peer Applications and Service Styles	474
28.2.2	A Referential View of Peer-to-Peer Interaction Styles	475
28.2.3	Business Models and Revenue Models	476
28.3	Revenue Models for P2P Business Application/Service Styles	477
28.3.1	Instant Messaging	477
28.3.2	Digital Content Sharing	480
28.3.3	Grid Computing	484
28.3.4	Collaboration	486
28.4	Discussion	487
29.	Peer-to-Peer Market Management	491
29.1	Requirements	491
29.1.1	Main Problems	492
29.1.2	Functional Requirements	493
29.1.3	Non-functional Requirements	493
29.2	Architecture	495
29.2.1	Market Model	495
29.2.2	Service Usage Model	497
29.2.3	Peer Model	497
29.2.4	Key Elements and Mechanisms	499
29.3	Case Studies	500
29.3.1	Peer-to-Peer Middleware	501
29.3.2	PeerMart: Peer-to-Peer Auctions	503
29.4	Conclusion and Outlook	507

30. A Peer-to-Peer Framework for Electronic Markets	509
30.1 Markets as Peer-to-Peer Systems	509
30.1.1 Service and Distribution Basics	510
30.1.2 SESAM Project Structure	512
30.2 A Service-Oriented Peer-to-Peer Architecture	513
30.2.1 Service Orientation	514
30.2.2 ServiceNets	515
30.2.3 Peer Architecture	517
30.3 Security, Robustness, and Privacy Challenges	519
30.3.1 Attack Classification/Threat Analysis	519
30.3.2 Peer-to-Peer-Related Challenges	520
30.3.3 Selected Issues	522
30.4 Summary	524

Part X. Advanced Issues

31. Security-Related Issues in Peer-to-Peer Networks	529
31.1 Introduction	529
31.2 Security Concerns on the Application Layer	529
31.2.1 File Sharing Applications	530
31.2.2 Data Backup Service	530
31.2.3 File Storage Service	531
31.3 Security Concerns on the Networking Layer	532
31.3.1 Invalid Lookup	532
31.3.2 Invalid Routing Update	533
31.3.3 Partition	533
31.3.4 Sybil Attack	534
31.3.5 Consideration of Implications of Topology	535
31.4 Security Concepts for Selected Systems	535
31.4.1 Groove	536
31.4.2 SixFour (6/4) Peer-to-Peer	540
31.4.3 Freenet	541
31.4.4 Further Peer-to-Peer Anonymizing Solutions	543
31.5 Conclusion	545
32. Accounting in Peer-to-Peer-Systems	547
32.1 The Purpose of Accounting	547
32.2 Why Is Not Accounting in Peer-to-Peer Straight Forward?	548
32.3 A Classification of P2P Accounting Schemes	549
32.3.1 Information Collection	549
32.3.2 Information Storage	551
32.4 Proposed Accounting Schemes	553
32.4.1 Plain Numbers-Based Systems	553
32.4.2 Receipt-Based Systems	555

32.4.3	Token-Based Systems	555
32.4.4	Proof of Work-Based Systems	556
32.5	Token-Based Accounting Scheme	556
32.5.1	Prerequisites	556
32.5.2	Overview	556
32.5.3	Token Structure	557
32.5.4	Token Aggregation	558
32.5.5	Check for Double Spending	559
32.5.6	Transactions	560
32.5.7	Trust & Security Considerations	561
32.5.8	Performance Analysis	564
32.5.9	Summary & Conclusions	566
33.	The PlanetLab Platform	567
33.1	Introduction and History	567
33.2	Architectural Principles	569
33.2.1	Application-Centric Interfaces	570
33.2.2	Distributed Virtualization	571
33.2.3	Unbundled Management	573
33.3	PlanetLab Methodology	574
33.3.1	Using PlanetLab	574
33.3.2	Reproducibility	575
33.3.3	Representivity	576
33.3.4	Quantitative Results	577
33.3.5	Qualitative Experience	578
33.4	Effects on the Internet	578
33.4.1	Many-to-Many Connections	579
33.4.2	Many Alternative Routes	579
33.4.3	Overlays and Traffic Correlation	580
33.5	Long-Term Goals	580
	Bibliography	583
	Index	623