

P

Peer-to-Peer

- 2 STEFAN SCHMID¹, ROGER WATTENHOFER²
 3 ¹Telekom Laboratories/TU Berlin, Berlin, Germany
 4 ²ETH Zürich, Zurich, Switzerland

Synonyms

6 Distributed hash table (DHT); Overlay network;
 7 Decentralization; Open distributed systems; Consistent
 8 hashing

Definition

10 The term *peer-to-peer* (p2p) is ambiguous, and is used
 11 in a variety of different contexts, such as:

- 12 • In popular media coverage, p2p is often synony-
 13 mous to software or protocols that allow users to
 14 “share” files (music, software, books, movies, etc.).
 15 p2p file sharing is very popular and a large fraction
 16 of the total Internet traffic is due to p2p.
- 17 • In academia, the term p2p is used mostly in two
 18 ways. A narrow view essentially defines p2p as
 19 the “theory behind file-sharing protocols.” In other
 20 words, how do Internet hosts need to be organized
 21 in order to deliver a search engine to find (share)
 22 content (files) efficiently? A popular term is “dis-
 23 tributed hash table” (DHT), a distributed data struc-
 24 ture that implements such a content search engine.
 25 A DHT should support at least a search (for a
 26 key) and an insert(key, object) operation. A DHT
 27 has many applications beyond file sharing, e.g., the
 28 Internet domain name system (DNS).
- 29 • A broader view generalizes p2p beyond file shar-
 30 ing: Indeed, there is a growing number of applica-
 31 tions operating outside the juridical gray area, e.g.,
 32 p2p Internet telephony à la Skype, p2p mass player
 33 games, p2p live audio&video streaming as in PPLive,
 34 StreamForge or Zattoo, or p2p social storage and
 35 cloud computing systems such as Wuala. Trying to

account for the new applications beyond file shar- 36
 ing, one might define p2p as a large-scale distributed 37
 system that operates without a central server bottle- 38
 neck. However, with this definition almost “every- 39
 thing decentralized” is p2p! 40

- From a different viewpoint, the term p2p may also 41
 be synonymous for privacy protection, as various 42
 p2p systems such as Freenet allow publishers of 43
 information to remain anonymous and uncensored. 44

In other words, there is no single well-fitting defini- 45
 tion of p2p, as some definitions in use today are even 46
 contradictory. In the following, an academic viewpoint 47
 is assumed (second and third definition above). 48

Discussion

49

The Paradigm

50

At the heart of p2p computing lies the idea that each 51
 network participant serves both as a producer (“server”) 52
 and consumer (“client”) of services. Depending on the 53
 application, the shared resources can be data (files), 54
 CPU power, disk storage, or network bandwidth. Often 55
 p2p systems have an open clientele, and do not rely on 56
 the availability of specific individual machines; rather 57
 they can deal with dynamic resources and do not exhibit 58
 single points of failure or bottlenecks. 59

Compared to centralized solutions, the p2p paradigm 60
 features a better scalability because the amount of 61
 resources grows with the network size, availability 62
 (avoiding a single point of failure), reliability, fair- 63
 ness, cooperation incentives, privacy, and security – 64
 just about everything researchers expect from a future 65
 Internet architecture. As such, it is not surprising that 66
 new “clean slate” Internet architecture proposals often 67
 revolve around p2p concepts. 68

One might naively assume that for instance scalabil- 69
 ity is not an issue in today’s Internet, as even most pop- 70
 ular web pages are generally highly available. However, 71
 this is not necessarily due to our well-designed Internet 72

73 architecture, but rather due to the help of so-called over-
74 lay networks: The Google Web site for instance man-
75 ages to respond so reliably and quickly because Google
76 maintains a large distributed infrastructure, essentially
77 a p2p system. Similarly, companies like Akamai sell “p2p
78 functionality” to their customers to make today’s user
79 experience possible in the first place. Quite possibly
80 today’s p2p applications are just testbeds for tomorrow’s
81 Internet architecture.

82 Implications

83 p2p networks are often highly dynamic in nature. While
84 traditional computer systems are typically based on
85 fixed infrastructures and are under a single adminis-
86 trative domain (e.g., owned and maintained by a single
87 company or corporation), the participating machines in
88 p2p networks are under the control of individual (and
89 to some extent: anonymous) users who can join and
90 leave at any time and concurrently. In p2p parlor, such
91 membership changes are called *churn*.

92 A second implication of the autonomy of the
93 machines in p2p networks is that the network consists
94 of different stakeholders. Users can have various reasons
95 for joining the network. For instance, an (anonymous)
96 user may not voluntarily contribute his or her band-
97 width, disk space, or CPU cycles to the system, but
98 prefer to *free ride*. This adds a socioeconomic aspect
99 to p2p computing. As the p2p paradigm relies on the
100 contributions of the participating machines, effective
101 incentive mechanisms have to be designed, which foster
102 cooperation and punish free riders.

103 Another source of inequality in p2p systems apart
104 from selfishness is *heterogeneity*: Due to the open mem-
105 bership, different machines run different operating sys-
106 tems, have different Internet connections, and so on.

107 Applications

108 The best-known representatives of p2p technology are
109 probably the numerous file-sharing applications such as
110 Napster, Gnutella, KaZaA, eMule, or BitTorrent. Also,
111 the Internet telephony tool Skype is very popular and
112 used by millions everyday. Zattoo, PPLive, and Stream-
113 Forge, among many others, use p2p principles to stream
114 video or audio content. The cloud computing service
115 Wuala offers free online storage by exploiting the par-
116 ticipants’ disks and Internet connections to improve

performance. Recently, the power and anonymity of
117 decentralized Internet working has gained the atten-
118 tion of operators of botnets in order to attack cer-
119 tain infrastructure components by a denial-of-service
120 attack. Finally, p2p technology is used for large-scale
121 computer games. 122

Architecture Variants

Several p2p architectures are known:

- Client/Server goes p2p: Even though Napster is
125 known to be the first p2p system (1999), by today’s
126 standards its architecture would not deserve the
127 label p2p anymore. Napster clients accessed a cen-
128 tral server that managed all the information of the
129 shared files, i.e., which file was to be found on
130 which client. Only the downloading process itself
131 was between clients (“peers”) directly, hence p2p.
132 In the early days of Napster the load of the server
133 was relatively small, so the simple Napster archite-
134 cture was sufficient. Over time, it turned out that the
135 server may become a bottleneck – and an attractive
136 target for an attack. Indeed, eventually a judge ruled
137 the server to be shut down (a “juridical denial of ser-
138 vice attack”). However, it remains to note that many
139 popular P2P networks today still include centralized
140 components, e.g., KaZaA or the eDonkey network
141 accessed by the eMule client. Also, the peer swarms
142 downloading the same file in the BitTorrent network
143 are organized by a so-called tracker whose function-
144 ality today is still centralized (although initiatives
145 exist to build distributed trackers). 146
- Unstructured p2p: The Gnutella protocol is the
147 antithesis of Napster, as it is a fully decentralized sys-
148 tem, with no single entity having a global picture.
149 Instead each peer connects to a random sample of
150 other peers, constantly changing the neighbors of
151 this virtual overlay network by exchanging neigh-
152 bors with neighbors of neighbors. (Any unstruc-
153 tured system also needs to solve the so-called
154 bootstrap *problem*, namely how to discover a first
155 neighbor in a decentralized manner. A popular solu-
156 tion is the use of well-known peer lists.) The fact
157 that users often turn off their clients once they
158 downloaded their content implies high levels of
159 churn (peers joining and leaving at high rates), and
160 hence selecting the right “random” neighbors is an
161

162 interesting research problem. The Achilles' heal of
 163 unstructured p2p architectures such as Gnutella is
 164 the cost of searching. A search request is typically
 165 flooded in the network and each search operation
 166 will cost m messages, m being the number of virtual
 167 edges in the architecture. In other words, such an
 168 unstructured p2p architecture will not scale. Indeed,
 169 when Napster was unplugged, Gnutella broke down
 170 as well soon afterward due to the inrush of former
 171 Napster users.

172 • Hybrid p2p: The synthesis of client/server archi-
 173 tectures such as Napster and unstructured archi-
 174 tectures such as Gnutella are hybrid architectures.
 175 Some powerful peers are promoted to so-called
 176 superpeers (or, similarly, trackers). The set of super-
 177 peers may change over time, and taking down
 178 a fraction of superpeers will not harm the sys-
 179 tem. Search requests are handled on the superpeer
 180 level, resulting in much less messages than in flat/
 181 homogeneous unstructured systems. Essentially, the
 182 superpeers together provide a more fault-tolerant
 183 version of the Napster server, as all regular peers
 184 connect to a superpeer. As of today, almost all pop-
 185 ular p2p systems have such a hybrid architecture,
 186 carefully trading off reliability and efficiency.

187 • Structured p2p: Inspired by the early success of
 188 Napster, the academic world started to look into
 189 the question of efficient file sharing. Indeed, even
 190 earlier, in 1997, Plaxton et al. [34] proposed a
 191 hypercubic architecture for p2p systems. This was a
 192 blueprint for many so-called structured p2p archi-
 193 tecture proposals, such as Chord [46], CAN [36],
 194 Pastry [37], Tapestry [50], Viceroy [26], Kadem-
 195 lia [27], Koorde [15], SkipGraph [3], and Skip-
 196 Net [11]. Maybe surprisingly, in practice, structured
 197 p2p architectures did not take off yet, apart from cer-
 198 tain exceptions such as the Kad architecture (from
 199 Kademlia [27]), which is accessible with the eMule
 200 client.

201 Scientific Origins

202 The scientific foundations of p2p computing were laid
 203 many years before the most simple "real" p2p systems
 204 like Napster emerged. As already mentioned, in 1997,
 205 a blueprint for structured systems has been proposed
 206 in [34]. Indeed, also the [34] paper was standing on the

shoulders of giants. Some of its eminent precursors are 207
 the following: 208


- Research on linear and consistent hashing, e.g., [16]. 209
- Research on locating shared objects, e.g., [4] or [5]. 210
- Research on so-called compact routing: The idea is 211
 to construct routing tables such that there is a trade- 212
 off between memory (size of routing tables) and 213
 stretch (quality of routes), e.g., [31] or [49]. 214
- ... and even earlier: hypercubic networks, see 215
 below! 216

Hypercubic Overlays and Consistent Hashing 217 218

Every application run on multiple machines needs a 219
 mechanism that allows the machines to exchange infor- 220
 mation. A naive solution is to store at each machine the 221
 domain name or IP address of every other machine. 222
 While this may work well for a small number of 223
 machines, large-scale distributed applications such as 224
 file sharing, grid computing, cloud computing, or data 225
 center networking systems need a different, more scal- 226
 able approach: instead of forming a clique (where every- 227
 body knows everybody else), each machine should 228
 only be required to know some small subset of other 229
 machines. This graph of knowledge can be seen as a 230
 logical network interconnecting the machines; it is also 231
 known as an *overlay network*. A prerequisite for an over- 232
 lay network to be useful is that it has good topological 233
 properties. Among the most important are small peer 234
 degree, small network diameter, robustness to churn, or 235
 absence of congestion bottlenecks. 236

The most basic network topologies used in practice 237
 are trees, rings, grids, or tori. Many other suggested net- 238
 works are simply combinations or derivatives of these. 239
 The advantage of trees is that the routing is very easy: for 240
 every source-destination pair there is only one possible 241
 path. However, the root of a tree can be a severe bottle- 242
 neck. An exception is a p2p streaming system where the 243
 single content provider forms the network root. How- 244
 ever, trees are also highly vulnerable, e.g., with respect 245
 to membership changes. 246

Essentially all state-of-the-art p2p networks today 247
 have some kind of hypercubic topology (e.g., Chord, 248
 Pastry, Kademlia). Hypercube graphs have many inter- 249
 esting properties, e.g., they allow for efficient routing: 250
 although each peer only needs to store a logarithmic 251



Peer-to-Peer. Fig. 1 A simplified p2p topology: a three-dimensional hypercube. Each peer has a three-bit identifier. For example, peer 110 is connected to the three peers 010, 100, 111 whose identifiers differ at exactly one position. In order to route a message from peer 110 to say peer 001, one bit is fixed after the other. One possible routing path is depicted in the figure: $110 \rightarrow 010 \rightarrow 000 \rightarrow 001$. An alternative path could be $110 \rightarrow 111 \rightarrow 101 \rightarrow 001$

252 number of other peers in the system (the peers' neigh-
 253 bors), by a simple routing scheme, a peer can reach each
 254 other peer in a logarithmic number of steps (or "hops").
 255 In a nutshell, this is achieved by assigning each peer a
 256 unique d -bit identifier. A peer is connected to all d peers
 257 that differ from its identifier at exactly one bit position.
 258 In the resulting hypercube network, routing is done by
 259 adjusting the bits in which the source and the destina-
 260 tion peers differ – one at a time (at most d many). Thus,
 261 if the source and the destination differ by k bits, there
 262 are $k!$ routes with k hops. Figure 1 gives an example.

263 Given a hypercubic topology, it is then simple to
 264 construct a distributed hash table (DHT): Assume there
 265 are $n = 2^d$ peers that are connected in a hypercube
 266 topology as described above. Now a globally known
 267 hash function f is used, mapping file names to long
 268 bit strings. Let f_d denote the first d bits (prefix) of
 269 the bitstring produced by f . If a peer is searching for
 270 file name X , it routes a request message $f(X)$ to peer
 271 $f_d(X)$. Clearly, peer $f_d(X)$ can only answer this request
 272 if all files with hash prefix $f_d(X)$ have been previously
 273 registered at peer $f_d(X)$.


274 There are some additional issues to be addressed
 275 in order to design a DHT from a hypercubic topol-
 276 ogy, in particular how to allow peers to join and leave
 277 without notice. To deal with churn the system needs
 278 some level of replication, i.e., a number of peers, which
 279 are responsible for each prefix such that failure of some

peers will not compromise the system. In addition, there
 are security and efficiency issues that can be addressed
 to improve the system.

280
 281
 282
 283 There are many hypercubic networks that are
 284 derived from the hypercube: among these are the but-
 285 tterfly, the cube-connected-cycles, the shuffle-exchange,
 286 and the de Bruijn graph. For example, the butter-
 287 fly graph is basically a "rolled out" hypercube (hence
 288 directly providing replication!) of constant degree.
 289 Another important class of hypercubic topologies are
 290 skip graphs [3, 11].

A simple, interesting way to design dynamic p2p
 systems is the *continuous-discrete approach* described
 by Naor and Wieder [29]. This approach is based on a
 "think continuously, act discretely" strategy, and can be
 used to design a variety of hypercubic topologies. The
 continuous-discrete approach gives a unified method
 for performing join/leave operations and for dealing
 with the scalability issue, thus separating it from the
 actual network. The idea is as follows: Let I be a
 Euclidean space, e.g., a (cyclic) one-dimensional space.
 Let G_c be a graph where the vertex set is the contin-
 uous set I . Each point in I is connected to some other
 points. The actual network then is a discretization of
 this continuous graph based on a dynamic decompo-
 sition of the underlying space I into cells where each
 "server" is responsible for a cell. Two cells are connected
 if they contain adjacent points in the continuous graph.
 Clearly, the partition of the space into cells should be
 maintained in a distributed manner. When a join oper-
 ation is performed, an existing cell splits, when a leave
 operation is performed two cells are merged into one.
 The task of designing a dynamic and scalable network
 follows these design rules: (1) Choose a proper con-
 tinuous graph G_c over the continuous space I . Design
 the algorithms in the continuous setting, which is often
 simpler (also in terms of analysis) than in the discrete
 case. (2) Find an efficient way to discretize the con-
 tinuous graph in a distributed manner, such that the
 algorithms designed for the continuous graph would
 perform well in the discrete graph. The discretization is
 done via a decomposition of I into the cells. If the cells
 that compose I are allowed to overlap, then the resulting
 graph would be fault tolerant.

320
 321
 322
 323
 324 To give an example, in order to build a dynamic *de*
 325 *Bruijn* network (a so-called Distance Halving DHT), a
 326 peer at position $x \in [0, 1)$ (in binary form $b_1b_2\dots$ such



Peer-to-Peer. Fig. 2 The continuous–discrete approach for the dynamic de Bruijn graph. Peers are indicated using circles, files using rectangles. In the continuous setting, the peer at position $x_i = 0.011010$ (in binary notation) is connected to positions $x_i/2$ and $(1+x_i)/2$. In the discrete setting, it is responsible for the cell (i.e., the connections and files that are mapped there) between positions x_i and x_{i+1}

327 that $x = \sum_{i=1}^{\infty} 2^{-b_i}$ connects to positions $l(x) := x/2 \in$
 328 $[0,1)$ and $r(x) := (1+x)/2 \in [0,1)$ in G_c (out-degree
 329 two per peer). Observe that if position x is written in
 330 binary form, then $l(x)$ effectively shifts in a “0” from the
 331 left and $r(x)$ shifts in a “1” from the left. Thus, routing
 332 is straightforward: based solely on the current position
 333 and the destination (without the overhead of maintain-
 334 ing routing tables), a message can be forwarded by a
 335 peer by fixing one bit per hop. The set of peers in the
 336 cyclic $[0,1)$ space then define the p2p network: Let x_i
 337 denote the position of the i^{th} peer (ordered in increas-
 338 ing order with respect to position). Peer i is responsible
 339 for the cell $[x_i, x_{i+1})$, computed in a modulo manner,
 340 i.e., this peer is responsible to store the data mapped to
 341 this cell plus for the establishment of the corresponding
 342 connections defined in G_c . Figure 2 gives an example.

343 Dealing with Churn

344 A distinguishing property of p2p systems are the fre-
 345 quent membership changes. Measuring the churn lev-
 346 els of existing p2p systems is challenging and one has
 347 to be careful when generalizing a given measurement
 348 to entire application classes (e.g., [10]). Nevertheless,
 349 several insightful measurement studies have been con-
 350 ducted. For instance, [9, 38] reported on the dynamic

nature of early p2p networks such as Napster and 351
 Gnutella, and [41] analyzed low-level data of a large 352
 Internet Service Provider (ISP) to estimate churn. Also 353
 the Kad DHT has been subject to measurement studies, 354
 and the reader is referred to the results in [47] and [43]. 355

It is widely believed that hypercubic structures are a 356
 good basis for churn-resilient p2p systems. As written 357
 earlier, a DHT is essentially a hypercubic structure with 358
 peers having identifiers such that they span the ID space 359
 of the objects to be stored. A simple approach to map the 360
 ID space onto the peers has already been described for 361
 the hypercube. To give another example, in the butterfly 362
 network, we may use its layers for replication, i.e., all 363
 peers with the same ID redundantly store the data of the 364
 same hash prefix. Other hypercubic DHTs can be more 365
 difficult to design, e.g., networks based on the pancake 366
 graph [19]. 367

For many well-known systems, theoretic analyses 368
 exist showing that the networks remain well-structured 369
 after some joins, leaves, or failures occur. In order to 370
 evaluate the robustness formally, metrics such as the 371
 network *expansion* (for deterministic failures) or the 372
span [6] (for randomized failures) are used. Unfortu- 373
 nately, the span is difficult to compute, and the span 374
 value is known only for the most simple topologies. 375

The continuous–discrete approach [29] already 376
 mentioned constitutes the basis of several dynamic sys- 377
 tems. For example, the SHELL system [40] is robust 378
 to certain attacks by connecting older or more reliable 379
 peers in a core network where access can be controlled; 380
 SHELL also allows to organize heterogeneous peers in 381
 an efficient topology. 382

Many systems proposed in the literature offer a 383
 high robustness in the average case, i.e., they provide 384
 probabilistic guarantees that hold with high probabilit- 385
 y. Robustness under attacks or worst-case dynamics 386
 is less well understood. In [19], a system is developed 387
 that achieves an optimal worst-case robustness in the 388
 sense that there is no alternative system that can toler- 389
 ate higher churn rates without disconnecting. The basic 390
 idea is to simulate a hypercube: each peer is part of a 391
 distinct hypercube *node*; each hypercube node consists of a 392
 logarithmic number of peers. Peers have connections to 393
 other peers of their hypercube node and to peers of the 394
 neighboring hypercube nodes. After a number of joins 395
 and leaves, some peers may have to change to another 396
 hypercube node such that up to constant factors, all 397

398 hypercube nodes have the same cardinality at all times.
 399 If the total number of peers grows or shrinks above or
 400 below a certain threshold, the dimension of the hyper-
 401 cube is increased or decreased by one, respectively. The
 402 balancing of peers among the nodes can be seen as a
 403 *dynamic token distribution problem* on the hypercube:
 404 Each node of a graph (hypercube) has a certain number
 405 of tokens, and the goal is to distribute the tokens along
 406 the edges of the graph such that all nodes end up with
 407 the same or almost the same number of tokens. Thus,
 408 the system builds on two basic components: (1) an algo-
 409 rithm, which performs the described dynamic token
 410 distribution and (2) an information aggregation algo-
 411 rithm, which is used to estimate the number of peers
 412 in the system and to adapt the hypercube's dimension
 413 accordingly. These techniques also work for alternative
 414 graphs, like pancake graphs [19].

415 An appealing notion of robustness is *topological self-*
 416 *stabilization*: A p2p topology is called self-stabilizing
 417 if it is guaranteed that from any weakly connected
 418 initial state (e.g., after an attack), it will quickly con-
 419 verge to a desirable network in the absence of fur-
 420 ther membership changes. In contrast to the worst-case
 421 churn considered in [19], self-stabilization focuses on
 422 the convergence time in periods without membership
 423 changes, but allows for general initial system states.
 424 While until recently, self-stabilizing algorithms with
 425 guaranteed runtime have only been known for sim-
 426 ple one-dimensional or two-dimensional linearization
 427 problems [14], recently a construction for a variation of
 428 skip graphs, namely SKIP+ graphs [13], has been pro-
 429 posed. Single joins and leaves in SKIP+ can be handled
 430 locally, and require logarithmic time and polylogarith-
 431 mic work only. However, there remains the important
 432 open question of how to provide degree guarantees
 433 during convergence from arbitrary states.

434 Fostering Cooperation

435 The appeal of p2p computing arises from the collab-
 436 oration of the system's constituent parts, the peers.
 437 If all the participating peers contribute some of their
 438 resources, highly scalable decentralized systems can
 439 be built. However, in reality, peers may act selfishly
 440 and strive for maximizing their own utility by ben-
 441 efitting from the system without contributing much

themselves [42]. Hence, the performance of a p2p sys- 442
 tem crucially depends on its capability of dealing with 443
 selfishness. 444

Already in 2000, Adar and Huberman [1] noticed 445
 that there exists a large fraction of free riders in the 446
 file-sharing network Gnutella. The problem of selfish 447
 behavior in p2p systems has been a hot topic in p2p 448
 research ever since, and many mechanisms to encour- 449
 age cooperation have been proposed [30]. Perhaps the 450
 simplest fairness mechanism is to directly incorporate 451
 contribution monitoring into the client software. For 452
 instance, in the file-sharing system KaZaA, the client 453
 records the contribution of its user. However, such a 454
 solution can simply be bypassed by implementing a dif- 455
 ferent client that hard-wires the contribution level to the 456
 maximum, as it was the case with *KaZaA Lite*. Inspired 457
 by real economies, some researchers have also proposed 458
 the introduction of some form of virtual money, which 459
 is used for the transactions. 460

BitTorrent has incorporated a fairness mechanism 461
 from the beginning and has hence been subject to 462
 intensive research (e.g., [21, 22, 35]). Although this 463
 mechanism has similarities to the well-known tit-for- 464
 tat scheme, the strategy employed in BitTorrent dis- 465
 tinguishes itself from the classic mechanism in many 466
 respects. For instance, it is possible for peers to obtain 467
 parts of a file "for free," i.e., without reciprocating. 468
 While this may be a useful property for bootstrapping 469
 newly joined peers, it has been shown that the Bit- 470
 Torrent mechanism can be exploited: the *BitThief* Bit- 471
 Torrent client [24] allows to download entire files fast 472
 without uploading any data. It has also been demon- 473
 strated in [24] that sharing communities are particularly 474
 vulnerable to such exploits. *BitThief* is not the only 475
 client cheating BitTorrent. Piatek et al. [32] presented 476
BitTyrant. *BitTyrant*'s strategy is to exploit the BitTor- 477
 rent protocol in order to maximize download rates. 478
 For instance, *BitTyrant* uses a smart neighbor selection 479
 strategy and connects to those peers with the best recip- 480
 rocation ratios. In contrast to *BitThief*, *BitTyrant* does 481
 not free ride. *BitTyrant* seeks to provide the minimal 482
 necessary contribution, and also increases the active 483
 neighbor set if this is beneficial to the download rate. 484
 The authors claim that their client provides a median 485
 70% performance gain in certain environments. 486

There can be many other forms of strategic behav- 487
 ior in open distributed systems. One subject that 488

489 has recently gained attention, especially by the game-
 490 theoretic research community, is *neighbor selection* in
 491 unstructured p2p networks (e.g., [28]). There may be
 492 several reasons for a peer to prefer connecting to some
 493 peers rather than others. For instance, a peer may want
 494 to connect to peers with high bandwidths, peers storing
 495 many interesting files, or peers having large degrees and
 496 hence provide quick access to many other peers. At the
 497 same time, a selfish peer itself may not be eager to store
 498 and maintain too many neighbors itself.

499 Current Trends and Outlook

500 One can argue that today, p2p computing is already
 501 a relatively mature (research) field; nevertheless, there
 502 are still many active discussions and developments, also
 503 in the context of the future Internet design. Moreover,
 504 there exists a discrepancy between the technology of the
 505 systems in use and what is actually known in theory.
 506 For example, the Kad network is still vulnerable to quite
 507 simple attacks [44].

508 If employed by the wrong people, the flexibility and
 509 robustness of p2p technology also constitutes a threat.
 510 Denial-of-service attacks are arguably one of the most
 511 cumbersome problems in today's Internet, and it is
 512 appealing to coordinate botnets in a p2p fashion. A
 513 DHT can be used by the bots, e.g., to download new
 514 instructions. For instance, it was estimated that in 2007,
 515 the DHT-based *Storm botnet* [20] ran on several million
 516 computers. Apart from mechanisms to detect or prevent
 517 attacks even before they take place, a smart redundancy
 518 management may improve availability during the attack
 519 itself (see, e.g., the Chameleon system [7]).

520 In terms of cooperation, there is a tension between
 521 the goal of providing incentive compatible mechanisms
 522 that exclude free riders and the goal of designing het-
 523 erogeneous p2p systems that also tolerate (and make
 524 use of!) weak participants. Moreover, in addition to
 525 design mechanisms dealing with pure selfishness, there
 526 is a trend toward p2p systems that are also resilient to
 527 malicious behavior (see, e.g., [23] or [39]).

528 Another active discussion regards the interface
 529 between p2p systems and ISPs. The large amount of
 530 p2p traffic raises the question of how ISPs should deal
 531 with p2p, e.g., by caching contents. p2p networks often
 532 employ inefficient overlay-to-ISP mappings as the logi-
 533 cal overlay network is typically not aware of the underly-
 534 ing "real" networks and constraints, and much overhead

can be avoided by improving the interface between p2p 535
 networks and ISPs, e.g., by an oracle [2]. For a criti- 536
 cal point of view on the subject, the reader is referred 537
 to [33]. 538

It seems that while a few years ago the lion's share 539
 of Internet traffic was due to p2p, the proportion seems 540
 to be declining [12] now. Especially web services and 541
 server-based solutions such as the popular YouTube and 542
 RapidShare are catching up. The measured data traces 543
 should be interpreted with care however, as they do not 544
 take into account what happens behind the scenes of 545
 big corporations. Indeed, it is believed that there is a 546
 paradigm shift in p2p computing: While p2p retreats 547
 (relatively to other applications) from public Internet 548
 traffic, today p2p technology plays a crucial role in the 549
 coordination and management of large data centers and 550
 server farms of corporations such as Akamai or Google. 551

Related Entries

►Cloud Computing	553
►Compact Routing	554
►Consistent Hashing	555
►Decentralization	556
►Distributed Hash Table (DHT)	557
►Grid Computing	558
►Hypercube	559
►Mechanism Design	560
►Open Distributed Systems	561
►Overlay	562
►Overlay Network	563
►Self-organization	564

Bibliographic Notes and Further Reading

Beyond the specific literature pointed to directly in 567
 the text, there are several recommendable introductory 568
 books on p2p computing. In particular, the reader is 569
 referred to the classic books [8, 45, 48] and two more 570
 recent issues [8, 17]. The theoretically more inclined 571
 reader may also be interested in [25], which provides an 572
 overview of compact routing solutions, and [18] which 573
 discusses trade-offs in local algorithms that achieve 574
 global goals based on local information only and with- 575
 out centralized entities whatsoever. Regarding the chal- 576
 lenges of distributed cooperation, the recent book [30] 577
 gives a thorough and up-to-date survey of current 578

579 (game-theoretic) trends, and also includes a chapter on
580 p2p specific questions.

581 Bibliography

- 582 1. Adar E, Huberman B (2000) Free riding on gnutella. *First Monday*
583 5(10):1–22
- 584 2. Aggarwal V, Feldmann A, Scheideler C (2007) Can ISPs and
585 p2p users cooperate for improved performance? *ACM Comput*
586 *Commun Rev* 37(3):29–40
- 587 3. Aspnes J, Shah G (2003) Skip graphs. In: *Proceedings of the 14th*
588 *annual ACM-SIAM symposium on discrete algorithms (SODA)*,
589 *Baltimore, 2003*
- 590 4. Awerbuch B, Peleg D (1990) Sparse partitions. In: *Proceedings of*
591 *the 31st annual symposium on foundations of computer science*
592 *(SFOCS), vol 2, pp 503–513, Washington, 1990*
- 593 5. Awerbuch B, Peleg D (1995) Online tracking of mobile users.
594 *JACM* 42(5):1021–1058
- 595 6. Bagchi A, Bhargava A, Chaudhary A, Eppstein D, Scheideler C
596 (2004) The effect of faults on network expansion. In: *Proceedings*
597 *of the 16th annual ACM symposium on parallelism in algorithms*
598 *and architectures (SPAA), Barcelona, 2004*
- 599 7. Baumgart M, Scheideler C, Schmid S. A DoSresilient information
600 system for dynamic data management. In: *Proceedings of the 21st*
601 *ACM symposium on parallelism in algorithms and architectures*
602 *(SPAA), Calgary, Alberta, 2009*
- 603 8. Buford J, Yu H, Lua EK (2008) P2P networking and applications.
604 Morgan Kaufmann, San Francisco, 2008
- 605 9. Gummadi K, Dunn R, Saroiu S, Gribble SD, Levy HM, Zahorjan J
606 (2003) Measurement, modeling, and analysis of a peer-to-peer
607 file-sharing workload. In: *Proceedings of the 19th ACM sympo-*
608 *sium on operating systems principles (SOSP), Bolton Landing,*
609 *2003*
- 610 10. Haeberlen A, Mislove A, Post A, Druschel P (2006) Fallacies
611 in evaluating decentralized systems. In: *Proceedings of the 5th*
612 *international workshop on peer-to-peer systems (IPTPS), Santa*
613 *Barbara, 2006*
- 614 11. Harvey NJA, Jones MB, Saroiu S, Theimer M, Wolman A. Skip-
615 net: a scalable overlay network with practical locality proper-
616 ties. In: *Proceedings of the 4th USENIX Symposium on Internet*
617 *Technologies and Systems (USITS), Seattle, 2003*
- 618 12. IPOQUE (2009) Internet study 2008/2009. [http://www.ipoque.](http://www.ipoque.com/resources/internet-studies/internet-study-2008-2009)
619 [com/resources/internet-studies/internet-study-2008-2009,](http://www.ipoque.com/resources/internet-studies/internet-study-2008-2009)
620 [pp 704–713 \(accessed on October 31, 2010\)](http://www.ipoque.com/resources/internet-studies/internet-study-2008-2009)
- 621 13. Jacob R, Richa A, Scheideler C, Schmid S, Täubig H (2009) A dis-
622 tributed polylogarithmic time algorithm for self-stabilizing skip
623 graphs. In: *Proceedings of the ACM symposium on principles of*
624 *distributed computing (PODC), New York, 2009*
- 625 14. Jacob R, Ritscher S, Scheideler C, Schmid S (2009) A self-
626 stabilizing and local delaunay graph construction. In: *Proceedings*
627 *of the 20th international symposium on algorithms and compu-*
628 *tation (ISAAC), Hawaii, 2009*
- 629 15. Kaashoek F, Karger DR (2003) Koorde: a simple degree-optimal
630 distributed hash table. In: *Proceedings of the international work-*
631 *shop on peer-to-peer systems (IPTPS), Berkeley, 2003*
16. Karger D, Lehman E, Leighton T, Panigrahy R, Levine M, Lewin D 632
(1997) Consistent hashing and random trees: distributed caching 633
protocols for relieving hot spots on the world wide web. In: *Pro-* 634
ceedings of the 29th ACM symposium on theory of computing 635
(STOC), New York, pp 654–663, 1997 636
17. Khan J, Wierzbicki A (2008) *Foundation of peer-to-peer comput-* 637
ing. Elsevier Computer Communication, 2008 638
18. Kuhn F, Moscibroda T, Wattenhofer R (2006) The price of being 639
near-sighted. In: *Proceedings of the 17th ACM-SIAM symposium* 640
on discrete algorithms (SODA), Miami, 2006 641
19. Kuhn F, Schmid S, Wattenhofer R (2010) Towards worstcase 642
churn resistant peer-to-peer systems. *J Distrib Comput (DIST)* 643
22(4):249–267 644
20. Larkin E (2007) Storm worm's virulence may change tactics. 645
British Computer Society (accessed on August 03, 2007) 646
21. Legout A, Urvoy-Keller G, Michiardi P (2006) Rarest first 647
and choke algorithms are enough. In: *Proceedings of the 6th* 648
ACM SIGCOMM conference on internet measurement (IMC), 649
pp 203–216, Rio de Janeiro, 2006 650
22. Levin D, LaCurts K, Spring N, Bhattacharjee B (2008) Bittor- 651
rent is an auction: analyzing and improving bittorrent's incentives. 652
SIGCOMM Comput Commun Rev 38(4):243–254 653
23. Li H, Clement A, Marchetti M, Kapritsos M, Robinson L, Alvisi L, 654
Dahlin M (2008) Flightpath: obedience vs choice in cooperative 655
services. In: *Proceedings of the symposium on operating systems* 656
design and implementation (OSDI), San Diego, 2008 657
24. Locher T, Moor P, Schmid S, Wattenhofer R (2006) Free riding in 658
bittorrent is cheap. In: *Proceedings of the 5th Workshop on Hot* 659
Topics in Networks (HotNets), Irvine, 2006 660
25. Malkhi D (2004) Locality-aware network solutions. Technical 661
Report, The Hebrew University of Jerusalem, HUJI-CSE-LTR- 662
2004-6 663
26. Malkhi D, Naor M, Ratajczak D (2002) Viceroy: a scalable 664
and dynamic emulation of the butterfly. In: *Proceedings of the* 665
21st annual symposium on principles of distributed computing 666
(PODC), Monterey, 2002 667
27. Maymounkov P, Mazières D (2002) Kademlia: a peer-to-peer 668
information system based on the xor metric. In: *Proceedings of* 669
the 1st international workshop on peer-to-peer systems (IPTPS), 670
Cambridge, 2002 671
28. Moscibroda T, Schmid S, Wattenhofer R (2006) On the topologies 672
formed by selfish peers. In: *Proceedings of the 25th annual sym-* 673
posium on principles of distributed computing (PODC), Denver, 674
2006 675
29. Naor M, Wieder U (2003) Novel architectures for p2p applica- 676
tions: the continuous-discrete approach. In: *Proceedings of the* 677
15th annual ACM symposium on parallel algorithms and archi- 678
tectures (SPAA), pp 50–59, San Diego, 2003 679
30. Nisan N, Roughgarden T, Tardos E, Vazirani VV (2007) *Algorith-* 680
mic game theory. Cambridge University Press, Cambridge 681
31. Peleg D, Upfal E (1988) A tradeoff between space and efficiency 682
for routing tables. In: *Proceedings of the 20th annual ACM sym-* 683
posium on theory of computing (STOC), pp 43–52, Chicago, 684
1988 685

- 686 32. Piatek M, Isdal T, Anderson T, Krishnamurthy A,
687 Venkataramani A (2007) Do incentives build robustness in
688 bittorrent? In: Proceedings of the 4th USENIX symposium on
689 networked systems design and implementation, Cambridge, 2007
- 690 33. Piatek M, Madhyastha HV, John JP, Krishnamurthy A,
691 Anderson T (2009) Pitfalls for ISP-friendly p2p design. In:
692 Proceedings of the hotnets, New York, 2009
- 693 34. Plaxton C, Rajaraman R, Richa AW (1997) Accessing nearby
694 copies of replicated objects in a distributed environment. In: Pro-
695 ceedings of the 9th ACM symposium on parallel algorithms and
696 architectures (SPAA), pp 311, 320, Newport, 1997
- 697 35. Qiu D, Srikant R (2004) Modeling and performance analysis
698 of bittorrent-like peer-to-peer networks. In: Proceedings of the
699 conference on applications, technologies, architectures, and pro-
700 tocols for computer communications (SIGCOMM), pp 367–378,
701 New York, 2004
- 702 36. Ratnasamy S, Francis P, Handley M, Karp R, Schenker S
703 (2001) A scalable content-addressable network. In: Proceedings
704 of the ACM SIG-COMM conference on applications, technolo-
705 gies, architectures, and protocols for computer communications,
706 pp 161–172, New York, 2001
- 707 37. Rowstron AIT, Druschel P (2001) Pastry: scalable, decentral-
708 ized object location, and routing for large-scale peer-to-peer
709 systems. In: Proceedings of the IFIP/ACM international confer-
710 ence on distributed systems platforms (middleware), pp 329–350,
711 Heibelberg, 2001
- 712 38. Saroiu S, Gummadi PK, Gribble SD (2002) A measurement study
713 of peer-to-peer file sharing systems. In: Proceedings of the multi-
714 media computing and networking (MMCN), San Jose, 2002
- 715 39. Scheideler C (2005) How to spread adversarial nodes?: rotate! In:
716 Proceedings of the 37th Annual ACM symposium on theory of
717 computing (STOC), pp 704–713, Baltimore, 2005
- 718 40. Scheideler C, Schmid S (2009) A distributed and oblivious heap.
719 In: Proceedings of the 36th international colloquium on automata,
720 languages and programming (ICALP), Rhodes, 2009
- 721 41. Sen S, Wang J (2004) Analyzing peer-to-peer traffic across large
722 networks. *IEEE/ACM Trans Netw* 12(2):219–232
- 723 42. Shneidman J, Parkes DC (2003) Rationality and self-interest in
724 peer to peer networks. In: Proceedings of the 2nd international
725 workshop on peer-to-peer systems (IPTPS), Berkeley, 2003
- 726 43. Steiner M, Biersack EW, Ennajary T (2007) Actively monitoring
727 peers in KAD. In: Proceedings of the 6th international workshop
728 on peer-to-peer systems (IPTPS), Bellevue, 2007
- 729 44. Steiner M, En-Najary T, Biersack EW (2007) Exploiting KAD:
730 possible uses and misuses. *Comput Commu Rev* 37(5):65–69
- 731 45. Steinmetz R, Wehrle K (2005) Peer-to-peer systems and applica-
732 tions. Springer, Heidelberg
- 733 46. Stoica I, Morris R, Karger D, Kaashoek F, Balakrishnan H (2001)
734 Chord: a scalable peer-to-peer lookup service for internet appli-
735 cations. In: Proceedings of the ACM SIGCOMM conference on
736 applications, technologies, architectures, and protocols for com-
737 puter communications, San Diego, 2001
- 738 47. Stutzbach D, Rejaie R (2006) Understanding churn in peer-to-
739 peer networks. In: Proceedings of the 6th internet measurement
740 conference (IMC), New York, 2006
48. Subramanian R, Goodman B (2005) Peer-to-peer computing: the
741 evolution of a disruptive technology. IGI, Hershey 742
49. Thorup M, Zwick U (2001) Compact routing schemes. In: Pro-
743 ceedings of the annual ACM symposium on parallel algorithms
744 and architectures (SPAA), pp 1–10, Crete, Greece, 2001 745
50. Zhao BY, Huang L, Stribling J, Rhea SC, Joseph AD, 746
Kubiatowicz J (2004) Tapestry: a resilient global-scale over-
747 lay for Service deployment. *IEEE journal on selected areas in*
748 *communucations* Vol. 22, No. 1 749