

- [1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in power-law networks. *Physical Review E*, 64:046135, 2001. [[bib](#) | [http](#)]

Many communication and social networks have power-law link distributions, containing a few nodes which have a very high degree and many with low degree. The high connectivity nodes play the important role of hubs in communication and networking, a fact which can be exploited when designing efficient search algorithms. We introduce a number of local search strategies which utilize high degree nodes in power-law graphs and which have costs which scale sub-linearly with the size of the graph. We also demonstrate the utility of these strategies on the Gnutella peer-to-peer network.

- [2] Avaki web site <http://www.avaki.com/>. [[bib](#) | [http](#)]
- [3] A. Bellissimo, P. Shenoy, and B. N. Levine. Exploring the use of BitTorrent as the basis for a large trace repository. Technical Report 04-41, Department of Computer Science, University of Massachusetts, June 2004. [[bib](#) | [.pdf](#)]
- [4] D. Brookshier, D. Govoni, N. Krishnan, and J. C. Soto. *JXTA: Java P2P Programming*. SAMS, Mar. 2002. [[bib](#)]
- [5] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker. Making Gnutella-like P2P systems scalable. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 407-418, 2003. [[bib](#) | [DOI](#)]

Napster pioneered the idea of peer-to-peer file sharing, and supported it with a centralized file search facility. Subsequent P2P systems like Gnutella adopted decentralized search algorithms. However, Gnutella's notoriously poor scaling led some to propose distributed hash table solutions to the wide-area file search problem. Contrary to that trend, we advocate retaining Gnutella's simplicity while proposing new mechanisms that greatly improve its scalability. Building upon prior research [1][27][36] we propose several modifications to Gnutella's design that dynamically adapt the overlay topology and the search algorithms in order to accommodate the natural heterogeneity present in most peer-to-peer systems. We test our design through simulations and the results show three to five orders of magnitude improvement in total system capacity. We also report on a prototype implementation and its deployment on a testbed.

- [6] Chord web site <http://www.pdos.lcs.mit.edu/chord/>. [[bib](#) | [http](#)]
- [7] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. *Lecture Notes in Computer Science*, 2009:46-66, July 2001. [[bib](#)]

We describe Freenet, an adaptive peer-to-peer network application that permits the publication, replication, and retrieval of data while protecting the anonymity of both authors and readers. Freenet operates as a network of identical nodes that collectively pool their storage space to store data files and cooperate to route requests to the most likely physical location of data. No broadcast search or centralized location index is employed. Files are referred to in a location-independent manner, and are dynamically replicated in locations near requestors and deleted from locations where there is no interest. It is infeasible to discover the true origin or destination of a file passing through the network, and difficult for a node operator to determine or be held responsible for the actual physical contents of her own node.

- [8] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 177-190, 2002. [[bib](#) | [DOI](#)]

The Peer-to-Peer (P2P) architectures that are most prevalent in today's Internet are decentralized and unstructured. Search is blind in that it is independent of the query and is thus not more effective than probing randomly chosen peers. One technique to improve the effectiveness of blind search is to proactively replicate data. We evaluate and compare different replication strategies and reveal interesting structure: Two very common but very different replication strategies-uniform and proportional-yield the same average performance on successful queries, and are in fact worse than

any replication strategy which lies between them. The optimal strategy lies between the two and can be achieved by simple distributed algorithms. These fundamental results offer a new understanding of replication and show that currently deployed replication strategies are far from optimal and that optimal replication is attainable by protocols that resemble existing ones in simplicity and operation.

- [9] C. Cranor, T. Johnson, O. Spataschek, and V. Shkapenyuk. Gigascope: a stream database for network applications. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pages 647-651, 2003. [[bib](#) | [DOI](#)]

We have developed *Gigascope*, a stream database for network applications including traffic analysis, intrusion detection, router configuration analysis, network research, network monitoring, and performance monitoring and debugging. Gigascope is undergoing installation at many sites within the AT&T network, including at OC48 routers, for detailed monitoring. In this paper we describe our motivation for and constraints in developing Gigascope, the Gigascope architecture and query language, and performance issues. We conclude with a discussion of stream database research problems we have found in our application.

- [10] Entropia web site <http://www.entropia.com/>. [[bib](#) | [http](#)]

- [11] F. L. Fessant, S. Handurukande, A.-M. Kermarrec, and L. Massoulie. Clustering in peer-to-peer file sharing workloads. In *Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, volume 3279 of *Lecture Notes in Computer Science*, pages 217-226. Springer, January 2005. [[bib](#) | [DOI](#)]

Peer-to-peer file sharing systems now generate a significant portion of Internet traffic. A good understanding of their workloads is crucial in order to improve their scalability, robustness and performance. Previous measurement studies on Kazaa and Gnutella were based on monitoring peer requests, and mostly concerned with peer and file availability and network traffic. In this paper, we take different measurements: instead of passively recording requests, we actively probe peers to get their cache contents information. This provides us with a map of contents, that we use to evaluate the degree of clustering in the system, and that could be exploited to improve significantly the search process.

- [12] Globus web site <http://www.globus.org/>. [[bib](#) | [http](#)]

- [13] Grid Forum web site <http://www.gridforum.org/>. [[bib](#) | [http](#)]

- [14] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, pages 314-329, 2003. [[bib](#) | [DOI](#)]

Peer-to-peer (P2P) file sharing accounts for an astonishing volume of current Internet traffic. This paper probes deeply into modern P2P file sharing systems and the forces that drive them. By doing so, we seek to increase our understanding of P2P file sharing workloads and their implications for future multimedia workloads. Our research uses a three-tiered approach. First, we analyze a 200-day trace of over 20 terabytes of Kazaa P2P traffic collected at the University of Washington. Second, we develop a model of multimedia workloads that lets us isolate, vary, and explore the impact of key system parameters. Our model, which we parameterize with statistics from our trace, lets us confirm various hypotheses about file-sharing behavior observed in the trace. Third, we explore the potential impact of locality-awareness in Kazaa. Our results reveal dramatic differences between P2P file sharing and Web traffic. For example, we show how the immutability of Kazaa's multimedia objects leads clients to fetch objects at most once; in contrast, a World-Wide Web client may fetch a popular page (e.g., CNN or Google) thousands of times. Moreover, we demonstrate that: (1) this "fetch-at-most-once" behavior causes the Kazaa popularity distribution to deviate substantially from Zipf curves we see for the Web, and (2) this deviation has significant implications for the performance of multimedia file-sharing systems. Unlike the Web, whose workload is driven by document change, we demonstrate that clients' fetch-at-most-once behavior, the creation of new objects, and the addition of new clients to the system are the primary forces that drive multimedia workloads such as Kazaa. We also show that there is substantial untapped locality in the Kazaa workload. Finally, we quantify the potential bandwidth savings that locality-aware P2P file-sharing architectures would achieve.

- [15] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, analysis, and modeling of bittorrent-like systems. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement (IMC)*, pages 35-48, 2005. [[bib](#) | [.html](#)]

Existing studies on BitTorrent systems are single-torrent based, while more than 85% of all peers participate in multiple torrents according to our trace analysis. In addition, these studies are not sufficiently insightful and accurate even for single-torrent models, due to some unrealistic assumptions. Our analysis of representative BitTorrent traffic provides several new findings regarding the limitations of BitTorrent systems: (1) Due to the exponentially decreasing peer arrival rate in reality, service availability in such systems becomes poor quickly, after which it is difficult for the file to be located and downloaded. (2) Client performance in the BitTorrent-like systems is unstable, and fluctuates widely with the peer population. (3) Existing systems could provide unfair services to peers, where peers with high downloading speed tend to download more and upload less. In this paper, we study these limitations on torrent evolution in realistic environments. Motivated by the analysis and modeling results, we further build a graph based multi-torrent model to study inter-torrent collaboration. Our model quantitatively provides strong motivation for inter-torrent collaboration instead of directly stimulating seeds to stay longer. We also discuss a system design to show the feasibility of multi-torrent collaboration.

- [16] A. Iamnitchi and I. Foster. On exploiting small-world usage patterns in file-sharing communities. In submission? [[bib](#)]
- [17] A. Iamnitchi, M. Ripeanu, and I. Foster. Small-world file-sharing communities. In *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, volume 2, pages 952-963, March 2004. [[bib](#) | [DOI](#)]

Web caches, content distribution networks, peer-to-peer file sharing networks, distributed file systems, and data grids all have in common that they involve a community of users who generate requests for shared data. In each case, overall system performance can be improved significantly if we can first identify and then exploit interesting structure within a community's access patterns. To this end, we propose a novel perspective on file sharing that considers the relationships that form among users based on the files in which they are interested. We propose a new structure that captures common user interests in data-the data-sharing graph-and justify its utility with studies on three data-distribution systems: a high-energy physics collaboration, the Web, and the Kazaa peer-to-peer network. We find small-world patterns in the data-sharing graphs of all three communities. We analyze these graphs and propose some probable causes for these emergent small-world patterns. The significance of small-world patterns is twofold: it provides a rigorous support to intuition and, perhaps most importantly, it suggests ways to design mechanisms that exploit these naturally emerging patterns.

- [18] A. Iamnitchi, M. Ripeanu, and I. T. Foster. Data-sharing relationships in the web. In *Poster at the 12th International World Wide Web Conference (WWW)*, May 2003. [[bib](#)]
- [19] Jini web site <http://www.sun.com/software/jini/>. [[bib](#) | [http](#)]
- [20] Project JXTA web site <http://www.jxta.org/>. [[bib](#) | [http](#)]
- [21] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, and M. Faloutsos. Is P2P dying or just hiding? In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, volume 3, pages 1532-1538, November/December 2004. [[bib](#) | [DOI](#)]

Recent reports in the popular media suggest a significant decrease in peer-to-peer (P2P) file-sharing traffic, attributed to the public's response to legal threats. Have we reached the end of the P2P revolution? In pursuit of legitimate data to verify this hypothesis, in this paper, we embark on a more accurate measurement effort of P2P traffic at the link level. In contrast to previous efforts, we introduce two novel elements in our methodology. First, we measure traffic of all known popular P2P protocols. Second, we go beyond the "known port" limitation by reverse engineering the protocols and identifying characteristic strings in the payload. We find that, if measured accurately, P2P traffic has never declined; indeed we have never seen the proportion of P2P traffic decrease over time (any change is an increase) in any of our data sources

- [22] R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In

This paper deals with the problem of spreading rumors in a distributed environment using randomized communication. In particular, we envisage the class of so-called epidemic algorithms which are commonly used for the lazy transmission of updates to distributed copies of a database. We introduce the random phone call model in order to investigate the possibilities and limits of this class of broadcasting algorithms. In this model, n players communicate in parallel communication rounds. In each round, each player calls a randomly selected communication partner. Whenever communication is established between two players, each one must decide which rumors to transmit. The major problem (arising due to the randomization) is that players do not know which rumors their communication partners have already received. In order to illustrate this problem, we will give a simple example of a commonly used algorithm in which each individual rumor is transmitted $\Theta(n \ln n)$ times. In this...

- [23] A. Klemm, C. Lindemann, M. K. Vernon, and O. P. Waldhorst. Characterizing the query behavior in peer-to-peer file sharing systems. In *Proceedings of the 4th ACM SIGCOMM conference on Internet Measurement (IMC)*, pages 55-67, October 2004. [[bib](#) | [DOI](#)]

This paper characterizes the query behavior of peers in a peer-to-peer (P2P) file sharing system. In contrast to previous work, which provides various aggregate workload statistics, we characterize peer behavior in a form that can be used for constructing representative synthetic workloads for evaluating new P2P system designs. In particular, the analysis exposes heterogeneous behavior that occurs on different days, in different geographical regions (i. e., Asia, Europe, and North America) or during different periods of the day. The workload measures include the fraction of connected sessions that are passive (i. e., issue no queries), the duration of such sessions, and for each active session, the number of queries issued, time until first query, query interarrival time, time after last query, and distribution of query popularity. Moreover, the key correlations in these workload measures are captured in the form of conditional distributions, such that the correlations can be accurately reproduced in a synthetic workload. The characterization is based on trace data gathered in the Gnutella P2P system over a period of 40 days. To characterize system-independent user behavior, we eliminate queries that are specific to the Gnutella system software, such as re-queries that are automatically issued by some client implementations to improve system responsiveness.

- [24] J. Kubiatoicz, D. Bindel, Y. Chen, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. Oceanstore: An architecture for global-scale persistent storage. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems Services*, Nov. 2000. [[bib](#) | [.html](#)]

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement of data. A prototype implementation is currently under development.

- [25] N. Leibowitz, M. Ripeanu, and A. Wierzbicki. Deconstructing the Kazaa network. In *Proceedings of the 3rd IEEE Workshop on Internet Applications (WIAPP)*, pages 112-120, June 2003. [[bib](#) | [http](#)]

Internet traffic is experiencing a shift from Web traffic to file swapping traffic. Today a significant part of Internet traffic is generated by peer-to-peer applications, mostly by the popular Kazaa application. Yet, to date, few studies analyze Kazaa traffic, thus leaving the bulk of Internet traffic in dark. We present a large-scale investigation of Kazaa traffic based on logs collected at a large Israeli ISP, which capture roughly a quarter of all traffic between Israel and US.

- [26] B. T. Loo, R. Huebsch, I. Stoica, and J. M. Hellerstein. The case for a hybrid P2P search infrastructure. In *Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, volume 3279 of *Lecture Notes in Computer Science*, pages 141-150. Springer, January 2005. [[bib](#) | [DOI](#)]

Popular P2P file-sharing systems like Gnutella and Kazaa use unstructured network designs. These networks typically adopt flooding-based search techniques to locate files. While flooding-based techniques are effective for locating highly replicated items, they are poorly suited for locating rare items. As an alternative, a wide variety of structured P2P networks such as distributed hash tables (DHTs) have been recently proposed. Structured networks can efficiently locate rare items, but they incur significantly higher overheads than unstructured P2P networks for popular files. Through extensive measurements of the Gnutella network from multiple vantage points, we argue for a hybrid search solution, where structured search techniques are used to index and locate rare items, and flooding techniques are used for locating highly replicated content. To illustrate, we present experimental results of a prototype implementation that runs at multiple sites on PlanetLab and participates live on the Gnutella network.

- [27] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th International Conference on Supercomputing (ICS)*, pages 84-95, 2002. [[bib](#) | [DOI](#)]

Decentralized and unstructured peer-to-peer networks such as Gnutella are attractive for certain applications because they require no centralized directories and no precise control over network topology or data placement. However, the flooding-based query algorithm used in Gnutella does not scale; each query generates a large amount of traffic and large systems quickly become overwhelmed by the query-induced load. This paper explores, through simulation, various alternatives to Gnutella's query algorithm, data replication strategy, and network topology. We propose a query algorithm based on multiple random walks that resolves queries almost as quickly as Gnutella's flooding method while reducing the network traffic by two orders of magnitude in many cases. We also present simulation results on a distributed replication strategy proposed in [8]. Finally, we find that among the various network topologies we consider, uniform random graphs yield the best performance.

- [28] M. Mannan and P. C. van Oorschot. On instant messaging worms, analysis and countermeasures. In *Proceedings of the ACM Workshop on Rapid Malcode (WORM)*, pages 2-11, 2005. [[bib](#) | [DOI](#)]

We provide a collection of minor results on the area of Instant Messaging (IM) worms, which has received relatively little attention in the formal literature. We review selected IM worms and summarize their main characteristics, motivating a brief overview of the network formed by IM contact lists, and a discussion of theoretical consequences of worms in such networks. Existing methods to restrict an IM worm epidemic are analyzed in terms of usability and effectiveness, leading to the suggestion of two minor variations to limit IM worm propagation. We believe these variations are more user-friendly and effective than existing published methods. We also provide brief results of a three and a half year user study of IM text messaging and file transfer frequency in a moderate-size public IM network - the largest such study to date - which is of independent interest, but also supports in part the preceding claim regarding user-friendliness.

- [29] E. P. Markatos. Tracing a large-scale peer to peer system: An hour in the life of Gnutella. In *Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid)*, pages 65-65, May 2002. [[bib](#) | [DOI](#)]

Peer-to-peer computing and networking, an emerging model of communication and computation, has recently started to gain significant acceptance. This model not only enables clients to take a more active role in the information dissemination process, but also may significantly increase the performance and reliability of the overall system, by eliminating the traditional notion of the "server" which could be a single point of failure, and a potential bottleneck. Although peer-to-peer systems enjoy significant and continually increasing popularity, we still do not have a clear understanding of the magnitude, the traffic patterns, and the potential performance bottlenecks of the recent peer-to-peer networks.

In this paper we study the traffic patterns of Gnutella, a popular large-scale peer-to-peer system, and show that traffic patterns are very bursty even over several time scales. We especially focus on the types of the queries submitted by Gnutella peers, and their associated replies. We show that the queries submitted exhibit significant amounts of locality, that is, queries tend to be frequently and repeatedly submitted. To capitalize on this locality, we propose simple Gnutella caching mechanisms

that cache query responses. Using trace-driven simulation we evaluate the effectiveness of Gnutella caching and show that it improves performance by as much as a factor of two.

[30] OceanStore web site <http://oceanstore.cs.berkeley.edu/>. [[bib](#) | [http](#)]

[31] Phex web site. [[bib](#) | [http](#)]

[32] Platform Computing web site <http://www.platform.com/>. [[bib](#) | [http](#)]

[33] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips. The Bittorrent P2P file-sharing system: Measurements and analysis. In *Proceedings of the 4th International Workshop on Peer-to-Peer Systems (IPTPS)*, volume 3640 of *Lecture Notes in Computer Science*, pages 205-216. Springer, 2005. [[bib](#) | [DOI](#)]

Of the many P2P file-sharing prototypes in existence, BitTorrent is one of the few that has managed to attract millions of users. BitTorrent relies on other (global) components for file search, employs a moderator system to ensure the integrity of file data, and uses a bartering technique for downloading in order to prevent users from freeriding. In this paper we present a measurement study of BitTorrent in which we focus on four issues, viz. availability, integrity, flashcrowd handling, and download performance. The purpose of this paper is to aid in the understanding of a real P2P system that apparently has the right mechanisms to attract a large user community, to provide measurement data that may be useful in modeling P2P systems, and to identify design issues in such systems.

[34] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 161-172, Aug. 2001. [[bib](#) | [.html](#)]

Hash tables-which map keys onto values-are an essential building block in modern software systems. We believe a similar functionality would be equally valuable to large distributed systems. In this paper, we introduce the concept of a Content-Addressable Network (CAN) as a distributed infrastructure that provides hash table-like functionality on Internet-like scales. The CAN is scalable, fault-tolerant and completely self-organizing, and we demonstrate its scalability, robustness and low-latency properties through simulation.

[35] S. Saroiu. *Measurement and Analysis of Internet Content Delivery Systems*. PhD thesis, University of Washington, December 2004. [[bib](#) | [.pdf](#)]

In recent years, the Internet has experienced an enormous increase in the use of specialized content delivery systems, such as peer-to-peer file-sharing systems (e.g., Kazaa, Gnutella, or Napster) and content delivery networks (e.g., Akamai). The sudden popularity of these systems has resulted in a flurry of research activity into novel peer-to-peer system designs. Because these systems: (1) are fully distributed, without any infrastructure that can be directly measured, (2) have novel distributed designs requiring new crawling techniques, and (3) use proprietary protocols, surprisingly little is known about the performance, behavior, and workloads of such systems in practice. Accordingly, much of the research into peer-to-peer networking is uninformed by the realities of deployed systems. This dissertation remedies this situation. We examine content delivery from the point of view of four content delivery systems: HTTP Web traffic, the Akamai content delivery network, and the Kazaa and Gnutella peer-to-peer file sharing networks. Our results (1) quantify the rapidly increasing importance of new content delivery systems, particularly peer-to-peer networks, and (2) characterize peer-to-peer systems both from an infrastructure and workload perspective. Overall, these results provide a new understanding of the behavior of the modern Internet and present a strong basis for the design of newer content delivery systems.

[36] S. Saroiu, P. Gummadi, and S. Gribble. A measurement study of peer-to-peer file sharing systems. In *Proceedings of Multimedia Computing and Networking (MMCN)*, January 2002. [[bib](#)]

[37] S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. *IEEE/ACM Transactions on Networking*, 12(2):219-232, 2004. [[bib](#) | [DOI](#)]

The use of peer-to-peer (P2P) applications is growing dramatically, particularly for sharing large video/audio files and software. In this paper, we analyze P2P traffic by measuring flow-level information collected at multiple border routers across a large ISP network, and report our investigation of three popular P2P systems—FastTrack, Gnutella, and Direct-Connect. We characterize the P2P traffic observed at a single ISP and its impact on the underlying network. We observe very skewed distribution in the traffic across the network at different levels of spatial aggregation (IP, prefix, AS). All three P2P systems exhibit significant dynamics at short time scale and particularly at the IP address level. Still, the fraction of P2P traffic contributed by each prefix is more stable than the corresponding distribution of either Web traffic or overall traffic. The high volume and good stability properties of P2P traffic suggests that the P2P workload is a good candidate for being managed via application-specific layer-3 traffic engineering in an ISP's network.

[38] C. Shirky. OpenP2P.com: What Is P2P ... And What Isn't, Nov. 2000. Article on web site. <http://www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html>. [[bib](#) | [.html](#)]

[39] R. D. Smith. Instant messaging as a scale-free network, May 2006. [[bib](#) | [http](#)]

The topology of an instant messaging system is described. Statistical measures of the network are given and compared with the statistics of a comparable random graph. The scale-free character of the network is examined and implications are given for the structure of social networks and instant messenger security.

[40] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. Technical Report TR-819, MIT, Mar. 2001. [[bib](#) | [.html](#)]

Efficiently determining the node that stores a data item in a distributed network is an important and challenging problem. This paper describes the motivation and design of the Chord system, a decentralized lookup service that stores key/value pairs for such networks. The Chord protocol takes as input an m -bit identifier (derived by hashing a higher-level application-specific key), and returns the node that stores the value corresponding to the key. Each Chord node is identified by an m -bit identifier and each node stores the key identifiers in the system closest to the node's identifier. Each node maintains an m -entry routing table that allows it to look up keys efficiently. Results from theoretical analysis, simulations, and experiments show that Chord is incrementally scalable, with insertion and lookup costs scaling logarithmically with the number of Chord nodes.

[41] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the ACM SIGCOMM '01 Conference*, pages 149-160, San Diego, California, Aug. 2001. [[bib](#) | [.pdf](#)]

A fundamental problem that confronts peer-to-peer applications is to efficiently locate the node that stores a particular data item. This paper presents Chord, a distributed lookup protocol that addresses this problem. Chord provides support for just one operation: given a key, it maps the key onto a node. Data location can be easily implemented on top of Chord by associating a key with each data item, and storing the key/data item pair at the node to which the key maps. Chord adapts efficiently as nodes join and leave the system, and can answer queries even if the system is continuously changing. Results from theoretical analysis, simulations, and experiments show that Chord is scalable, with communication cost and the state maintained by each node scaling logarithmically with the number of Chord nodes.

[42] I. Stoica, R. Morris, D. Liben-Novell, D. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. Technical report, MIT, Jan. 2002. [[bib](#) | [.ps](#)]

A fundamental problem that confronts peer-to-peer applications is to efficiently locate the node that stores a particular data item. This paper presents Chord, a distributed lookup protocol that addresses this problem. Chord provides support for just one operation: given a key, it maps the key onto a node. Data location can be easily implemented on top of Chord by associating a key with each data item, and storing the key/data item pair at the node to which the key maps. Chord adapts efficiently as nodes join and leave the system, and can answer queries even if the system is continuously

changing. Results from theoretical analysis and simulations show that Chord is scalable, with communication cost and the state maintained by each node scaling logarithmically with the number of Chord nodes.

- [43] R. Thommes and M. Coates. Epidemiological modelling of peer-to-peer viruses and pollution. In *Proceedings of the IEEE INFOCOM*, April 2006. [[bib](#) | [.pdf](#)]
- [44] K. Tutschku. A measurement-based traffic profile of the eDonkey filesharing service. In *Proceedings of the 5th International Workshop on Passive and Active Network Measurement (PAM)*, volume 3015 of *Lecture Notes in Computer Science*, pages 12-21. Springer, April 2004. [[bib](#) | [DOI](#)]

Peer-to-peer file sharing applications have evolved to one of the major traffic sources in the Internet. In particular, the eDonkey file sharing system and its derivatives are causing high amounts of traffic volume in today's networks. The eDonkey system is typically used for exchanging very large files like audio/video CDs or even DVD images. In this report we provide a measurement based traffic profile of the eDonkey service. Furthermore, we discuss how this type of service increases the "mice and elephants" phenomenon in the Internet traffic characteristics.

- [45] May 1997. URN Syntax. Internet Engineering Task Force. RFC 2141R.
<http://ietf.org/rfc/rfc2141.txt>. [[bib](#) | [.txt](#)]
- [46] B. J. Wilson. *JXTA*. New Riders, June 2002. [[bib](#) | [http](#)]

This file was generated by [bibtex2html](#) 1.96.