
Volume 31, issue 3, 25 February 2008 ISSN 0140-3664
ELSEVIER
computer
communications
Available online at
ScienceDirect
www.sciencedirect.com
Special issue:
Disruptive Networking with Peer-to-Peer Systems
Guest Editors: Javed I. Khan
Adam Wierzbicki

Available online at www.sciencedirect.com

ScienceDirect

Computer Communications 31 (2008) 419–422

computer communications

www.elsevier.com/locate/comcom

Introduction

# Guest editors' introduction: Disruptive networking with peer-to-peer systems

Peer-to-peer computing is the disruptive technology of our time. It is disruptive in many senses – even for networking. It has become experimentation ground for fundamentally new modes of communication. Historically, the success of influential communication technologies seems to be founded on their ability to reduce barriers of communication from their predecessor. The optical telegraph of Claude Chappe (1792) [1] and subsequent telegraph age drastically reduced the barrier of *one-to-one long distance* communication compared to the horse and carriage mode of the time. Some 80 years later telephone (1876) network took away the last-mile latency and made communication instantaneous. Telephone is a 1-to-1 mode. Though in principle it allows 1-to-*n* party communication but then the cost barrier creeps up in order of *n*. The advent of television broadcast (1925) flattened the barrier. To be precise it took away the barrier of receiving – though the cost barrier of sending remained sky high. Nevertheless, it firmly established itself as the dominant technology for 1-to-*n* unidirectional communication. These examples illustrate that the barriers to adoption of a communication technology are often not only technical, but economic or social in nature. The significance of the Internet is that it promised the holy grail – a technology for low barrier *n*-to-*n* modes – where no technology of previous generations could go. This goal is hard to achieve even with Internet; bare-bone TCP/IP does not quite suffice. Yet, we believe that using the Internet, low barrier many-to-many communication can be achieved. Interestingly, the widespread adoption of the Internet's early communication models has led to their apparent ossification, making it hard to advance the Internet's abilities to a new level. What was needed was a way to overcome the Internet's limitations without disrupting the operations of its successful legacy applications, and without incurring a prohibitive cost.

During the last two decades a rich set of modes has been constructed. Email, web, groups, social networks, recent blogging are each an example of such constructions over the Internet. One of the first, *email*, is a 1-to-1 communication technology (if not extended by e-mail lists, when it becomes 1-to-*n*). A page posted in a *web server* also offered a low barrier 1-to-*n* mode. Yet, while the client–server model offers an easy way to communicate content from a server to many clients, it introduces bottlenecks and a single point of failure that limit its scalability. Perhaps more significant is the limitation of the client–server model when it comes to the ability of publishing content by clients. The current development of Web 2.0 [2] shows how important this ability is to current Internet users, who are not satisfied to be consumers, but take the role of producers of information. The client–server model allows such a development, but again at an increased cost. But to many it really could not be the ultimate *n*-to-*n*.

In this context, peer-to-peer systems entered into the picture and ventured to shatter this publishing barrier. All started with a brick-and-mortar architectural goal where each participant can upload as easily as download. This has turned peer-to-peer research into a bold experimentation about new modes of low barrier *n*-to-*n* communication. Routing issues of *n*-to-*n* communication are investigated in depth in various works under the publish/subscribe overlays design banner.

Perhaps one of the interesting innovations in contemporary networking is BitTorrent. It practically demonstrated that the more popular is the resource, the faster it can be downloaded in a *n*-to-*n* world defying the commonsense law of hotspots. What is actually demonstrated is the fundamental advantage of *concurrent communication* – despite the resistance to it in classical TCP/IP world. Indeed more such disruptions are possible by shattering few more assumptions we take for granted in classical communication science and engineering. Here are few things to ponder.

## 1. Pathway concurrency

In classical networking bytes travel via a singular path from source to sink- a well-behaved path is more like a series points or a tree, not a graph-wide flow. Is such restriction really needed? It seems a session may use multiple parallel flows resulting in many benefits. We can call it *pathway concurrency*. Unlike the early days of networking connectivity is no longer scarce – most places now have concurrent reachability all the way to the end-hosts. The

0140-3664/$ - see front matter 
doi:10.1016/j.comcom.2007.10.001

Internet is now turning more mesh-like – than original skinny sparse graph. Redundant links are abundant. Even laptops are now embellished with multiple adapters. So why not use parallel communication more? Parallelism can be exploited ingenuously towards countless advantages ranging from increased capacity and reliability, low latency, even sound security. Assumption of ample redundancy in connectivity and parallelism can potentially flatten routing. Route table explosion has become the choking point in classical architecture. With bandwidth multiplicative technology like CWDM – there will be more link abundance and more fingers will be pointer at routers and switches. We have to look into something fundamentally new for next level of scale up. The P2P community seems to be already experimenting with some of those new modes. P2P architectural processes start from a dense graph topology assumption. From foundational point of view can DHT research (context dependent forwarding, distinction between short range/long range forwarding) be viewed as the harbinger of such new routing over a quasi-mesh?

## 2. Link concurrency

Another accepted norm in classical network's notion of path is the need of an abridged chain of *simultaneously available links* from source to sink. All it's links must be simultaneously up for a path to exist. Store-and-forwarding is a cornerstone piece of TCP/IP architecture – yet the concept of storage is severely underdeveloped. Without storage, the restriction of link simultaneity can be relaxed at best up to split seconds. More or less the case of real store-and-forward communication is treated as an alien requirement in this architecture. Any case of communication that involves link in-concurrency has to seek help from outside of network's architecture. Yet the general notion of communication found in natural systems, seldom imposes such a restriction. In the general notion of 'connectivity' – requirement of link concurrency is much looser. The treatments of concurrent vs. non-concurrent-link pathways are more holistic and seamless. Matured treatment of non-concurrency of links is fundamentally important especially when the communication is $n$-to-$n$, the system is large, and the bodies involved are autonomous. In the P2P model, the treatment of churn, self-organization of topology and overlays, and seamless treatment of store-and-forward nature of communication contributes right here. The recent development of research on Delay-Tolerant Networks (DTNs) [3] can be seen as a natural consequence of P2P research that has demonstrated that percolation (the phenomenon of communication that is resistant to long delays, outages, or transient lack of connectivity) is indeed possible in real-life communication systems. In an environment with link in-concurrency apparently only two fundamental strategies are possible. One way of handling it is to make network intelligence fluid where most nodes are to be prepared to play all roles. The other way is message storage – where information can wait out the period of link unavailability.

Interestingly, the starting tenet of P2P paradigm is the former and the meta-feature of most of the services over P2P is the latter.

P2P also upsets the cost perception in architecture. In classical architecture the general purpose nodes are at the edge and the core is made up of special purpose machines such as routers and switches. Support for more generalized communication stemming from either type of concurrency requires painful change at the core to classical networks. Hence, there is a big resistance in classical networking to experiment with any fundamentally new mode of communication. In P2P architecture the distinction between edge and core vanishes. Both the edge and the core here have general purpose machines. Thus, in P2P paradigm there is no extra cost for experimenting with all these new modes of communications. There are several obvious consequences of this elimination. Core is now empowered to do storage, intelligent rerouting or almost anything. But also fine-grained autonomy seeps inside core, resulting in churn.

This volume of the special issue showcases a set of 16 papers grouped under pub–sub overlay, multicast overlay, and synchronous overlay, and organization and architecture. The volume is by no means exhaustive – but rather a snap-shot of the innovative communication issues and techniques researchers in peer-to-peer architecture is pursuing. Out of more than hundred submissions, these papers have been selected for the importance of the issue, creativity in solution approach and their grounding on formal analytic techniques.

## 3. Pub–sub overlays

The routing issues involved in $n$-to-$n$ communication are being investigated in P2P community under the banner of pub–sub overlays. Here, a group of n can communicate by publishing and subscribing about a topic set t. The information is routed via these topic set. Alberto Mozo and Salvachúa present a hybrid unstructured–structured pub/sub network where a topical tag attached to published objects is stored in intermediate nodes using fast and compact bloom filter and then the search process finds right objects by their tag with in $O(\log(n))$ node hop bound. Cutting, Quigley, and Landfeldt show how hot topics can be further handled by exploiting parallelism. They suggest dividing it between matched groups and show how to optimally sized the groups for load balance.

## 4. Dynamic multicast

Some $n$-to-$n$ communication forms can be decomposed into a set of dynamic 1-to-$n$ multicasts where any node can become the source anytime. A corresponding distribution tree to other $n-1$ nodes needs to be built dynamically without any significant cost barrier. Three papers in this issue delve into the topic of *on-the-fky multicasting*. Buford, Brown, and Kolberg present an interesting scheme for multi-destination multicast messaging based on another scheme of parallel communication. They show how it achieves near lower bounds of Chuang-Sirbu multicast scaling law. Lua and his team present a scheme for forming a stable tree under churn. They suggest a two-layer peer-to-peer multicasting network built using an innovative 'geometric map' derived from internet distances. Pompili, Scoglio, and Lopez show how to satisfy users varying QoS demands by maintaining an optimum virtual shared trees among group members.

## 5. Synchronous overlays

There is a strong push in P2P community to open up the sharing of multimedia documents – such as streaming video or video on demand (VoD). Essentially it means experimenting with finer flow and data concurrency constraints inside the *n*-to-*n* communication sessions. Do, Hua, and Tantaoui present a P2P architecture aiming at VoD where each peer only caches a moving window of the ongoing stream and continually forwards the rest along a distribution tree of caches. A concept called 'generation' is introduced to manage subscribers joining at different times. The design of the ad hoc distribution tree for media streaming is non-trivial and can be optimized on many criterion. Guo, Suh, Kurose, and Towsley present a VoD technique for scalability. Byun and Yoo show how to optimally place the gateway/caches placement when a set of peers might be willing to act as a video distribution gateway (perhaps for a fee) with linear programming and Lagrangian Relaxation. Akbari, Rabiee, and Ghanbari show how to allocate bandwidth in an existing overlay spanning tree so that all peers are optimally satisfied for their respective upper QoS requirement. Haridasan and vanRenesse investigate the security in streaming and present a protocol that they claim can avert many malicious behaviors at the end level.

## 6. Organization and architecture

The peer-to-peer architectures are classically divided into two forms. Structured DHT systems offer performance bounds while the unstructured overlays show flexibility. The later is very important for sustained growth and scale. Design goal of peer-to-peer systems is to find a balance between the two. It seems as the peer-to-peer networks grow, some form of hierarchical organization will be inevitable to strike this balance. Churn or dynamic availability of peers is yet another factor that further complicates P2P system. Zoels, Despotovic, and Kellerer show how to optimally divide a network into two levels and ensure optimum performance, when nodes have varying degree of availability. Yu and Li address how to optimally organize and place Trackers in a BitTorrent like P2P using fluid dynamic like analysis and propose clustered BitTorrent. Kersch, Szabó, Cheng, Jean, and Galis presents a two prong connection maintenance mechanism – a proactive self-stabilizing strategy for short range connections and a probabilistic strategy for long range connection to build a stable network despite churn. It retains the flexibility as well the DHT like logarithmic bound on performance. Carchiolo, Malgeri, Mangioni, and Nicosia present another interesting approach to connect the peers in an efficient overlay using the 'semantics' of the contents. All routing and management are also done based on strictly local information. The results are quite intriguing. This resulting overlay seems to capture many benefits of small-world networks. They show strong clustering, community structure and small average distances leading to efficient search. Due to inherent node autonomy in P2P paradigm there is no completeness. This makes every P2P application a compromise between completeness of scope and completion time. This gives rise to the interesting new problem of cascaded wait-time management in asynchronous 1-to-*n* autonomous-body communication. The next paper from editor's group (Khan and Haque) present schemes for cascaded nodes to decide how long to wait or not to wait before it can forward/return its current harvest. It reports dramatic improvement in search time. Many of us have often wondered at the intuitive similarity between mobile and ad hoc networks and peer-to-peer systems. If churn is not enough to make things complex, how about try building a peer-to-peer system out of wondering nodes? Zhan and Schiller takes on this ultimate problem and show how indeed such a system might be practically built without paying overwhelming cost for the continuous restructuring of physical network.

Peer-to-peer paradigm currently is a development over the TCP/IP substrate. It is offering critical experimentation ground for new concepts in communication. It enables development and deployment of protocols – which the TCP/IP constricted networking layers did not permit. It is quite possible that some of the striking findings will flow back into lower layers eventually for their intrinsic strength. Distribution of intelligence, scale, and individual autonomy are the three hallmarks of peer-to-peer communication. These are also the hallmark attribute of the ultimate Internet – the all-encompassing conduit of social communication. For the time being, peer-to-peer paradigm seems to be the foremost experimentation ground to build the technology for that. We hope not only the P2P researchers, but also researchers in other layers of networking will benefit from the innovations occurring right now in peer-to-peer paradigm.

## References

[1] Gerard J. Holzmann, Bjorn Pehrson, The Early History of Data Networks, ISBN: 978-0-8186-6782-4, December 1994, Wiley-IEEE Computer Society Press.

[2] Tim O'Reilly, What Is Web 2.0, Design Patterns and Business Models for the Next Generation of Software, September, 2005, (Available from: <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>, Last retrieved 9/20/2007).

[3] Kevin Fall, A delay-tolerant network architecture for challenged internets, in: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, August 25–29, 2003, Karlsruhe, Germany (doi:10.1145/863955.863960).

Javed I. Khan *
*Networking and Media Communications,*
*Laboratories Department of Computer Science,*
*Kent State University, USA*
*E-mail address:* javed@kent.edu

Adam Wierzbicki
*Polish-Japanese Institute of Information Technology,*
*ul. Koszykowa 86, 02-008 Warsaw, Poland*
*E-mail address:* adamw@pjwstk.edu.pl

Available online 6 October 2007

* Corresponding author.