

FairE9: Fair File Distribution over Mesh-Only Peer-to-Peer

Eyal Zohar
Computer Science Department
The Open University of Israel
Email: eyalzo@gmail.com

Anat Lerner
Computer Science Department
The Open University of Israel
Email: anat@openu.ac.il

Abstract—Peer-to-peer (P2P) networks are in the spotlight due to the wide-spreading file-sharing applications. Many file-distributing algorithms have been suggested and implemented. The various solutions need to cope with a heterogeneous and unstable environment, where peers can arrive and depart at a high rate (churn). Sometimes cooperation cannot be assumed. These issues make the structured attitude less practical. Even some of the algorithms that are considered as unstructured try to maintain long-term parent-child relationships. Existing unstructured (mesh-only) algorithms for file-distribution work well on the Internet on the average. But some of the participating peers may suffer from a slow start or high latency because of the randomness of the peer and piece selection for upload and download.

In this paper we propose a fair unstructured system for file-distribution from a single source, with no central authority. The proposed protocol is fair both with respect to load balancing and with respect to the latency in each peer. It is based on a novel weights-algorithm that helps peers to determine what piece to ask from which peer, in a manner that increases their chance to get served. In this way it also lowers the overhead. The proposed algorithm welcomes newcomers while being resilient to churn, being resilient to free-riders, and adaptive to heterogeneous bandwidth.

Index Terms—Communication systems, Protocols.

I. INTRODUCTION

The peer-to-peer (P2P) system is a promising platform for many forthcoming distributed systems that deal with content delivery. The magic of P2P systems is based on the ability of a client to share previously downloaded data, thus helping numerous servers with their distribution. Broadcasting servers can deal with a large client-base by having the clients form a P2P overlay. With such an overlay, web sites and multimedia servers can cope with a sudden increase in service demand (flash-crowds) [1], [2].

P2P overlays can be classified into two main classes: mesh-only and non-mesh. The non-mesh solutions are mostly based on some (hidden) hierarchical structures, with long-term parent-child relationships. Initial effort is spent on the establishment of a structure, which in turn allows an efficient streaming of the information along the paths.

SplitStream [3] is a tree-based multicast system for a heterogeneous environment. It splits the video stream into stripes using MDC (Multiple Description Coding). With MDC, a peer can watch the video as soon as it receives at least one stripe, but to improve its quality it needs more stripes. SplitStream

distributes each stripe using a separate distribution tree. Many existing solutions follow the idea of SplitStream to build multiple trees/paths to distribute the forwarding load among all peers. Each algorithm may focus on a different property. For example, Orchard [4] deals with free-riding, Chunkyspread [5] adapts to heterogeneous capacity, LagOver [6] considers individual latency constraints, iGridMedia [7] has guaranteed delay, and Climber [8] and [9] are incentives-based.

These algorithms often lose their main advantage in the Internet environment, where peers arrive and leave (churn) frequently, and where some are unwilling or unable to share in return (free-riding). Broken paths are not easily detected, and when these are repaired they cost with some more overhead and a higher latency.

Mesh-only solutions divide files into small equal pieces to allow peers to upload soon after they start download. The peers are given a high level of freedom, as they can contact any peer, and download from it any piece the other peer owns and is willing to upload. Central entities, if they exist, do not fully control the process but merely help peers to locate each other. Famous representatives of this family are the existing file-sharing networks, like BitTorrent [10], Gnutella [11] and eMule [12]. These operate well in a heterogeneous bandwidth environment, and sometimes encourage peers to share. But their flexibility and openness come with some level of chaos which results in random performance and encourages peers to be aggressive by contacting many siblings. The distributed nature and direct-reciprocity help free-riders to bypass defense mechanisms by whitewashing (switching identity) and/or downloading from complete-sources (peers that completed their download) [13]. Mechanisms that are intended to filter out free-riders may also hurt potentially-cooperative newcomers [14].

In this paper we suggest a novel approach to the problem, which puts some order in a mesh-only solution. Each peer has an inherent individual piece-order that can be easily derived by all other peers from a commonly known function. The combination of the two peers' orders imposes weights to each request, helping to estimate the probability for a transaction before the requests are actually sent. Altogether this algorithm enhances fairness and lowers the overhead of the process.

This paper is organized as follows. Some background is given in Section II. The proposed algorithm is presented in

Section III. In Section IV we show experimental evaluations. We present several ideas for future work in Section V and conclude in Section VI.

II. BACKGROUND

FairE9 takes a novel approach toward the file distribution problem. To our knowledge, there are no other mesh-only P2P networks that are designed to cope with flash-crowd. In this section we give a brief overview of some systems that will later be referenced and compared to.

A. BitTorrent

Direct reciprocity P2P networks are based on the concept of direct reward given by a downloader to a peer that uploaded to it.

BitTorrent is a pioneer of mesh-only P2P file-sharing, which can be regarded as having a direct-reciprocity mechanism with an immediate reward. The implemented mechanism is roughly based on TFT (Tit-for-Tat) [15]. Studies show effective upload bandwidth utilization [16]; however, it is also far from optimal for distribution from a single source to multiple receivers [17]. In addition, its peer selection policy induces free-riding [18], its upload bandwidth distribution is suboptimal [19], and measurements show that when the number of seeds is over 60%, free-riders download faster than others [13].

The distribution process is made up of many peers running a software client that implements the original BitTorrent wire protocol. Numerous software clients exist, each has its own flavor: from a basic open-source, to sophisticated clients that use their proprietary extensions. The incentive mechanism is not a part of the wire protocol, but is more related to internal decision taking. As with the wire protocol, each software client may take a slightly different approach.

Basically, each peer holds a fixed number of upload slots, hopefully allocated to cooperative and beneficial partners. Peers normally measure their download speed in these slots, in order to choke (send a specific "choke" message) or even disconnect the less productive ones, and look for better partners. In order for cooperation to get started, it also does an "optimistic unchoke" from time to time, meaning it gives a new partner an upload slot for a short time to see if it cooperates. This behavior is intended for newcomers to get their time-limited chance to get their first pieces, but it also helps free-riders to download.

B. eMule

eMule file sharing system was suggested earlier than BitTorrent. It is considered less efficient than BitTorrent, and it is less popular worldwide today. It implements a long-term credit system, in which peer a gains a remote credit when uploading to peer b , and loses some credit when downloading from it. If peer a later needs anything that b shares, it will move faster in b 's queue and get better service according to its credit there.

eMule's main drawback is long latency and large overhead. Each peer has a single incoming-requests queue that is common to all the files it shares. Peers are encouraged to contact

many potential sources and to send multiple messages just to keep their place in the remote queues until they get a chance to download one piece (around 10 MB). Sometimes it takes several days before a peer gets a significant part of a downloaded file.

III. FAIRE9 SCHEME

A. The Model

A group of peers $N = \{1, \dots, n\}$ is interested in a file¹ $f \in F$. The file is divided into m pieces, indexed with serial numbers $\{1, \dots, m\}$. Initially the file is held exclusively by the initial source $s \in N$. The goal is to distribute all the pieces to all the peers in N . When a peer completes downloading all the pieces of the file, it is called a *seeder*.

B. The Idea

Given a file, each peer maintains a strict piece order, determined only by its network address, i.e., its IP address. This piece order can also be independently computed based on the peer's IP by every other peer. A source, given multiple requests by multiple peers, determines its serving order according to its own piece order and the requesting peers piece order it computes locally. The algorithm uses the IP address as a parameter to a hash function and obtains an ID, enabling the permutation function to distribute the pieces uniformly over the permutation space. Hence, the piece order list determines the priority of the client per request.

A piece order list is created by a permutation function, g , that maps a triple of a peer, a file, and a piece number into the index in the peer's piece-order. The permutation creates a prioritized piece request order, such that if, for example, the third piece is placed first, the algorithm motivates the peer to request it first. The permutation function g converts a number into a permutation in a linear time [20], [21]. g is a bijection (one-to-one and onto) function, and it is known to all peers in N :

$$g : [N, F, \{1, \dots, m\}] \rightarrow \{1, \dots, m\}$$

Each peer repeatedly sends requests to other peers, asking to download a specific missing piece. Sources get these requests, and consider providing the needed piece. Each request is granted a calculated *weight*, such that the requests with the *lower* weights get higher priority and are served first. Weights are based on piece position on both sides. When peer con asks for piece i of file f from peer src , both peers measure the request's weight using the formula:

$$weight(con, src, f, i) = \begin{cases} g(con, f, i) + g(src, f, i) & \text{if } src \neq s \\ g(con, f, i) & \text{if } src = s \end{cases}$$

In the case of collisions in the permutation function that may result in a collision at a source, a momentary decision process at the source is reduced to a random one, resulting

¹FairE9 supports the distribution of multiple files in parallel, while each is distributed in a separate domain, like BitTorrent does.

with a winner and a loser. The loser is then first in line for the next grant.

1) *Peer Identifier*: The permutation function makes use of the peer identifier. A reasonable peer identifier mechanism would be the IP address of a peer, converted to a serial integer (up to 2^{16}). Other alternatives may be considered, as long as they follow the basic required properties of such an identifier. The identity should not be freely selected by the peer itself, although non-frequent switches are allowed. FairE9 does not save any credits, therefore multiple peers with the same identifier may exist, just not too often; groups of around $1/m$ of the peers can share the same identifier with no harm.

2) *Locating Peers*: FairE9 does not limit itself to one solution, but it can use any combination of one of the many operative solutions already existing in current networks; for example BitTorrent trackers, BitTorrent DHT, eMule servers, eMuleKademlia, etc. For the ease of description, from this point on, we assume that a BitTorrent trackers mechanism is in use.

C. The Random Protocol

In this section we describe a fully random algorithm that can be considered as the ancestor of several mesh P2P networks, like eMule, BitTorrent and FairE9.

Procedure 1 MainRandom

1. Once: *Bootstrap()*
 2. Forever, in background: *PeerExchange()*
 3. Forever, in background: *AcceptRequests()*
 4. Until seed, in background: *SendRequests()*
 5. Forever: *Upload()*
-

Procedure 2 Bootstrap

1. Get initial peers from external source
 2. Add peers to *known_peers*
 3. Add peers to *free_peers*
-

Procedure 3 PeerExchange

1. Get random peer from *known_peers*
 2. Ask peer for its *known_peers*
 3. **for all** returned peers, while not exceed *known_peers* **do**
 4. **if** returned is new to *known_peers* **then**
 5. Add returned to *known_peers*
 6. Add returned to *free_peers*
 7. **end if**
 8. **end for**
-

Procedure 4 AcceptRequests

1. Wait for incoming requests
 2. **if** not exceed *up_pending* **then**
 3. Add request for random piece to *up_pending*
 4. **end if**
-

Procedure 5 SendRequests

1. Cleanup expired *down_pending*
 2. Remove disconnected *known_peers/free_peers*
 3. **if** not exceed *down_pending* **then**
 4. Pop random peer from *free_peers*
 5. Send request for random piece to the peer
 6. Add the request to *down_pending*
 7. **end if**
-

Procedure 6 Upload

1. **if** *up_pending* is not empty and *up_slots* is not full **then**
 2. Pop random peer from *up_pending*
 3. **if** there are pieces you have and peer does not **then**
 4. Pick random piece of these pieces
 5. Occupy a slot in *up_slots*
 6. Upload the piece to the peer in background
 7. When complete, release a slot in *up_slots*
 8. **end if**
 9. **end if**
-

D. The FairE9 Protocol

FairE9 suggests a different approach to sending requests and to uploading:

Procedure 7 SendRequests'

1. Cleanup expired *down_pending*, and abandon the peers that were attached to the removed requests
 2. Remove disconnected *known_peers/free_peers*
 3. **if** not exceed *down_pending* **then**
 4. $target \leftarrow \max(2, down_pending / \text{missing pieces})$
 5. Get the lightest missing piece that has less than *target* pending requests
 6. Pop from *free_peers* the peer with the lightest weight for that piece
 7. Send the request for the piece to the peer
 8. Add the request to *down_pending*
 9. **end if**
-

Procedure 8 Upload'

1. **if** *up_pending* is not empty and *up_slots* is not full **then**
 2. Pop the lightest peer and piece from *up_pending*
 3. **if** you have that piece **then**
 4. Occupy a slot in *up_slots*
 5. Upload the piece to the peer in background
 6. When complete, release a slot in *up_slots*
 7. **end if**
 8. **end if**
-

E. FairE9 Protocol Behavior

Under the FairE9 protocol, peers have a prioritized download requests order, and a clear view of preferred peers from whom there is a high probability of being served. This

mechanism has several fundamental desirable properties and outcomes:

- 1) Newcomers are welcome in the sense that new peers will get their first pieces very fast.
- 2) Total overhead is reduced as the possibility of being served is increased.
- 3) Total load is fairly shared as peers will usually receive requests for their low-order pieces.
- 4) Requesting peers temporarily abandon peers that have not uploaded to them fast enough (expired *down_pending*), in order to cope with free-riders and slow peers.
- 5) Peers that are granted their first requests, are in need of pieces of higher order (higher index number) than pieces requested by slower or newer peers, and hence are less likely to be serviced first. Thus, our scheme maintains an on-the-fly order and fairness between all peers in the steady state.

IV. EXPERIMENTAL EVALUATION

In this section we present simulation results of experiments designed to evaluate the performance of FairE9 and compare them to existing algorithms. We ran large-scale experiments on a proprietary network simulator that implements several different systems in an almost real environment. The results support our claims that several aspects of fairness and overhead can be improved by the new system.

A. Simulation Setup

We developed a proprietary simulator in order to evaluate specific behaviors of the proposed system, while neutralizing mechanisms that are not in the scope of this paper, such as the peers acquiring mechanism.

All the simulated systems are evaluated using the same fixed parameters, unless noted otherwise: 10,000 peers download a 250MB file, with 200 pieces, from a single source; each peer, including the single source, has a 3Mbps download and 3Mbps upload bandwidths, divided into 500Kbps slots; each peer has 5 initial random known peers (previously joined), and performs a single peer exchange with one of the peers it knows every 40 seconds to get up to 50 new peers, until it reaches 500 known peers; maximum number of pending requests is set to 100, with timeout of 200 seconds per request.

To simulate eMule, we created an eMule-like mode that resembles Random, but uses eMule credits mechanism. Peers waiting in the queue get higher priority if they have uploaded to the same peer before. Each uploaded piece gives a boost equal to a 60 seconds wait in the queue.

BitTorrent does not have a strict united incentives mechanism because many software clients exist today, each taking a slightly different approach. Our BitTorrent-like mode implements the basic ideas presented by the original software client [22]; each peer regularly checks the download-to-upload ratio on each connection to keep the connections with the high ratios alive, and closes those with the lower ratios. It also periodically gives a chance to new peers for a limited time.

B. Flash Crowd and Fairness

We tested an extreme case of flash crowd, when the single source deals with all the newcomers simultaneously from the very beginning. This is a common situation with live events or timed site launches. The goal is to keep both the average latency (efficiency) and the maximum and standard deviation latencies (fairness) low.

This is a difficult assignment for mesh P2P that deals at start with peers that know only some of the other peers; none of them have a history or anything to give in return. We wish to avoid unfair situations, where some peers download several pieces while others starve for their first piece.

In Table I we compare Random, eMule and FairE9, with respect to the average, the standard deviation and the maximum latencies. The simulation results show that FairE9 is both more efficient and more fair than Random, eMule and BitTorrent.

To explain the differences, a closer look at the process is needed. Fig. 1 presents the percentage of slower peers that were at least 20 pieces (10% of the file) behind the average. Fig. 2 does the same for 40 pieces (20%) behind the average. The diagrams show the undesired impact of BitTorrent's TFT and eMule's credit system, where a relatively small group of peers is starved. These are unlucky peers that get their first pieces late; therefore they have lower credit and slow continuation, even in comparison to Random. FairE9 on the other hand helps peers that started late to catch up quickly.

TABLE I
FLASH CROWD RESULTS

Protocol	Avg. Latency	Latency Stdev	Max Latency
Random	795.4 sec	35.6	1,100 sec
eMule	812.8 sec	38.8	1,260 sec
BitTorrent	819.2 sec	47.6	1,300 sec
FairE9	786.4 sec	19.8	880 sec

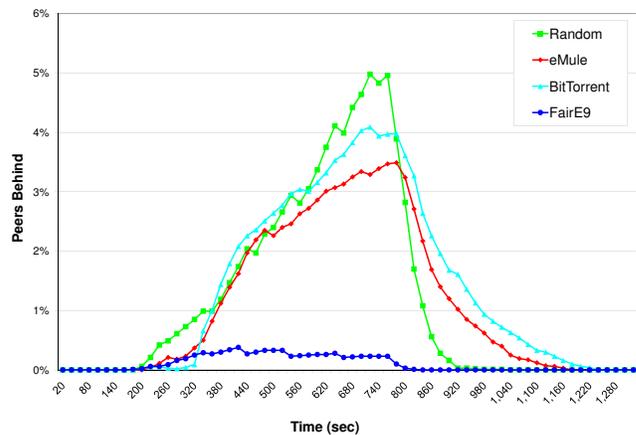


Fig. 1. Peers 10% behind the average

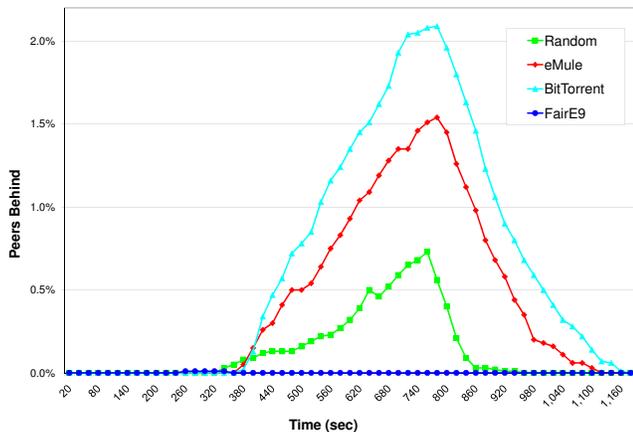


Fig. 2. Peers 20% behind the average

C. Load Balance and Fairness

Another aspect of fairness is related to the load carried by each peer. We have measured the number of uploads performed by each peer. The results in Table II show that FairE9 keeps the system relatively balanced, probably because it turns newcomers into "almost seeders" very fast, as they get their first pieces early, and these first pieces are also the pieces often requested from them by others.

BitTorrent demonstrates the worst results, probably because some of the peers are unlucky with their first pieces, and TFT only makes the situation worse. Not only do peers not get a chance to upload much, they also need to download mostly from seeders that do not ask for anything in return. Altogether it creates an unbalanced load.

D. Moderateness and Message Overhead

Typical mesh P2P encourages peers in practice to contact many peers, as the chance to download grows with the number of requests sent. What stops peers from sending unlimited number of requests are costs and constraints related to communications and CPU usage.

Consider for example a peer that sends i requests. Denote by p the probability to be served by at least one of the requested peers. Using the Random algorithm, if this peer doubles the number of requests, the probability to be served will grow to $2p - p^2$, as the probability to download from the peers of the second group is equal to the probability to download

TABLE II
UPLOADS DEVIATION

Protocol	Uploads Stdev.	Max Uploads
Random	29.52	288
eMule	26.01 sec	339
BitTorrent	31.94 sec	358
FairE9	20.95 sec	253

from the peers of the first group. With FairE9, the requests are ordered in a descending order of the probability to be served. Therefore, if we denote by q the probability to download from the second group, then $q \ll p$, and so doubling the number of requests will be less profitable.

To evaluate the motivation of a single peer in practice, we defined a peer *aggressiveness factor* which is a multiplier of the number of maximum pending requests (the base is 100). Table III shows what happens when the crowd has an aggressive factor of 2.0 while a tested group of 10 peers remains with a factor of 1.0. With respect to the reduced latency of the crowd and the added latency of the group, it is clear that Random is the worst. One could expect that the moderate group would not perform so well in eMule and BitTorrent because their somewhat random service encourages peers to be aggressive. But surprisingly the moderate group performs well in both networks because in this special experiment there are hardly any seeders during the process which makes the direct-reciprocity more dominant and random contacts less effective. The reason that FairE9 is influenced by the factor is because the aggressive peers have an advantage when downloading their last pieces.

Message overhead has some performance penalty, therefore in real life there is a trade-off between the reduced latency and the additional overhead. Therefore, we also measured the average number of messages per download sent by the peers. The results are presented in Table IV. Intuitively, as FairE9 sends messages for specific pieces while the others ask for any piece, the number of messages in FairE9 could be the highest. But, in the others, the peers have to exchange inventory messages to figure out if and what one can send to the other, so the final results are that FairE9 has a lower message overhead in this test case. BitTorrent has the highest overhead, probably because the unlucky newcomers need to send many messages until they get a chance to download. eMule is a little better, probably because peers mostly get a chance to download even without a credit, just by waiting long enough for their turn.

E. Welcoming Newcomers

A common problem with incentives-based systems is the way they treat the newcomers who cannot contribute anything or cannot express or prove their willingness to share.

To emphasize the effect of the approach taken by FairE9, we simulated a crowd that joins the distribution gradually; 8,000 peers join at time 0, and 200 more join every minute, up to

TABLE III
AGGRESSIVE CROWD VS. MODERATE GROUP

Protocol	Crowd avg latency reduction (1.0→2.0)	Group avg latency gain (2.0→1.0)
Random	18.2 sec	154 sec
eMule	0.6 sec	10 sec
BitTorrent	2.4 sec	6 sec
FairE9	12.2 sec	30 sec

TABLE IV
MESSAGE OVERHEAD

Protocol	Messages per download
Random	5.15
eMule	5.82
BitTorrent	6.53
FairE9	4.82

10,000 peers. For each group of 200 peers we measured the average latency and absolute completion time (join time plus latency). We found that the absolute completion time correlates with the join time on all the protocols and groups. But the latency results are more complex, as presented in Fig. 3.

Newcomers in eMule and BitTorrent usually pay for their late join by having a long latency. An exception to this is when they join the distribution in a progressed stage, with a high number of seeders (or almost seeders).

We also compared the delay for receiving the first piece by the newcomers, as shown in Fig. 4. The samples are taken one minute after a group of 200 peers joins, just before another group joins. It is clear that FairE9 and Random help newcomers to download, as almost all the 200 peers download at least one piece within the first minute that they join. eMule and BitTorrent incentive mechanisms treat newcomers as inferior peers, therefore they have difficulties getting their first piece.

F. Resilience to Seeders Churn

The problem with peers that leave unstructured systems is that they might take with them rare pieces of the file or just take with them a significant part of it while the overall progress is just at the beginning.

We have tested the case of peers that leave soon after they become seeders. Fig. 5 presents the portion of non-seeders still in the system, while seeders leave with probability 0.3 every 20 seconds. The results correlate with Fig. 2, because eMule's

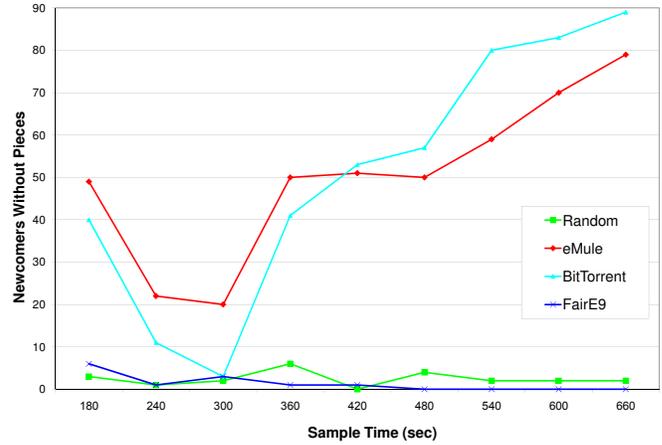


Fig. 4. Sampled newcomers

starving peers also suffer from the phenomenon explored here, while FairE9 is completely resilient even to a high churn rate.

G. Incentives to Share

Any comparison that involves eMule and BitTorrent is incomplete without an evaluation of the incentives to share, as this is their main concern and focus. To evaluate the incentives to share, we set a different limited upload bandwidth to several groups of 50 peers each. Fig. 6 shows the latency penalty paid by each group.

It is clear that eMule and BitTorrent impose incentives to upload. But it should be noted that we also measured the number of uploads performed by each group and saw that it correlates with the upload bandwidth presented here. So with these two protocols peers that have only pieces that others do not need, or a limited upload bandwidth, also suffer from long latency. Another important point is that even free-riders in BitTorrent complete their download eventually, paying less than additional 40% in latency. So the incentive to share

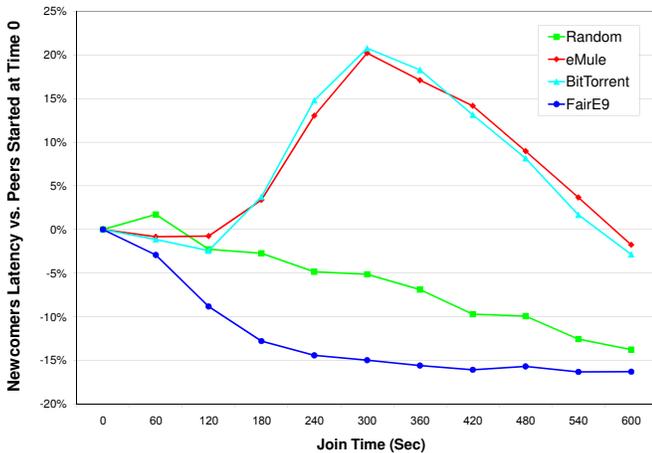


Fig. 3. Newcomers latency

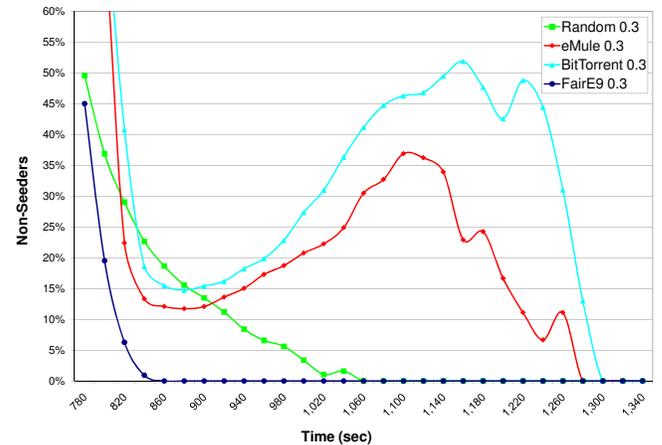


Fig. 5. Seeders churn, with leaving probability 0.3

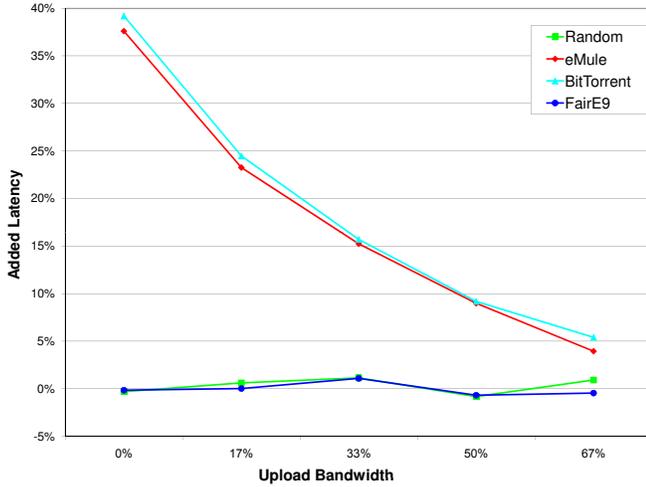


Fig. 6. Limited upload bandwidth

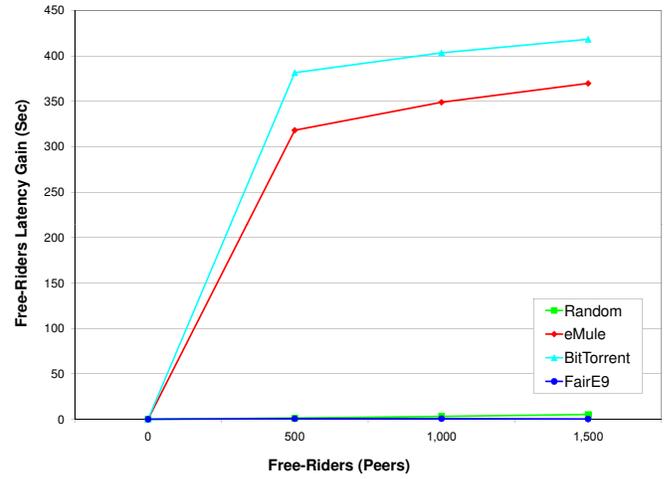


Fig. 8. Free-riders compared to sharing peers

depends on the trade-off between latency and dedicated upload bandwidth.

H. Free-Riders

Even though free-riders are not FairE9’s main concern, we find it necessary to compare their effect on FairE9 to their effect on eMule and BitTorrent.

In Fig. 7 we compare the effect on the average latency when some of the peers free-ride on the various algorithms. In Fig. 8 we focus on the latency gain of the free-riders in comparison to the sharing peers.

FairE9 takes a completely different approach to free-riders than eMule and BitTorrent. Instead of a direct punishment to the free-riding peer, a weak incentive to cooperate is offered, as cooperative behavior slightly decreases the overall and self latencies.

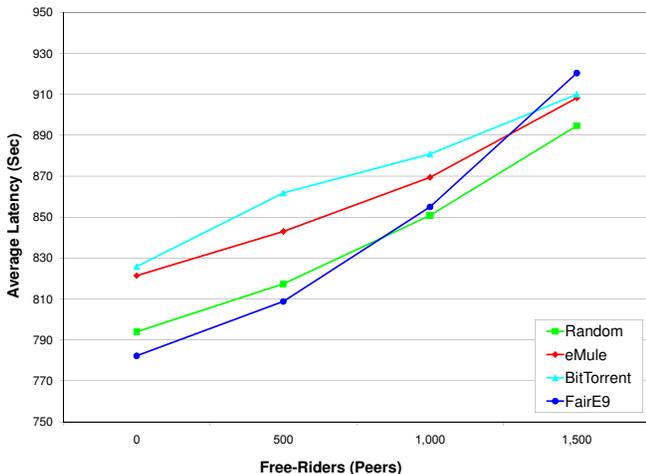


Fig. 7. Average latency in the presence of free-riders

I. Reduce Hard Disk Drive Reads

Sometimes a peer’s upload bottleneck is the memory reading speed. This is due to the hard disk drive access time that derives from its head movements. A broadband peer that shares many files made of small pieces, may encounter such a problem. One of the keys to the solution is to make each peer focus on a small number of distinct pieces to upload. This can not only reduce head movements but it also enables the use of fast cache memory.

Table V presents the results of a measured number of distinct pieces uploaded by each peer. Since peers start to share from the moment they have only one piece, even in the random protocols, peers upload less than half of the pieces. Random is a little better than eMule, because the latter sometimes couples peers that exchange distinct pieces on purpose. eMule is a little better than BitTorrent because of a similar reason, as BitTorrent’s TFT is a stricter decoupling mechanism.

FairE9 deals with this phenomenon by having each peer become in practice a major source for only a small number of distinct pieces. More than that, the weight system helps to focus on 1-2 pieces, as illustrated in Fig. 9. According to it a peer could serve 55% of 1,000 files from 250MB of fast memory.

V. FUTURE WORK

FairE9 can be enhanced to be incentives-based and even directly confront free-riders. Since peers are expected to

TABLE V
DISTINCT UPLOADED PIECES PER PROTOCOL

Protocol	Distinct Pieces Avg.
Random	91.83
eMule	93.71
BitTorrent	94.39
FairE9	37.58

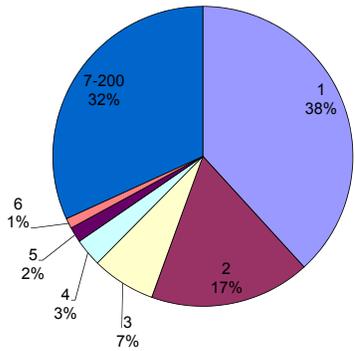


Fig. 9. FairE9 uploads by piece position

download according to piece-order, the system is ready for direct-reciprocity mechanism. One option is that a peer that wishes to download will be forced to upload simultaneously a lower order piece. This constraint is expected to add latency to the average case, with even more latency to less cooperative peers. Such a scheme does not require the use of long-term memory or the presence of a central authority.

The presented weight formula gives equal importance to the index of the piece on both sides. An unbalanced formula can change the properties of the protocol. We expect that a formula that gives a higher weight to the index on the downloading side will give a boost to newcomers and improve the latency. The trade-off would probably be additional overhead.

We used a balanced permutation function. In some cases it may be more suitable to use an unbalanced function. For example, if a video preview is required it would be better to use a function that tends to put the first file pieces in low indexes. This will distribute the first pieces faster than the other pieces.

VI. CONCLUSIONS

We presented the design and evaluation of FairE9, a fair system for file-distribution from a single source, with no central authority. Each peer has an inherent individual piece-order. The combination of the two peers' orders imposes weights on each request. Serving priority is determined by these weights, making the entire process more deterministic and systematic.

As other unstructured P2P networks, it copes well with churn, free-riders, heterogeneous environments and peer failures. In comparison to other unstructured P2P networks, FairE9 also handles flash crowd, balances the upload efforts, welcomes newcomers and reduces the message overhead.

ACKNOWLEDGMENTS

The authors wish to thank Zeev Nutov for helping with the permutations algorithm; and Ran Tavory for his comments.

REFERENCES

- [1] J. Jung, B. Krishnamurthy, and M. Rabinovich, "Flash crowds and denial of service attacks: characterization and implications for cdns and web sites," in *WWW '02: Proceedings of the 11th international conference on World Wide Web*. New York, NY, USA: ACM, 2002, pp. 293–304.
- [2] R. Kumar and K. W. Ross, "Peer-assisted file distribution: The minimum distribution time," in *Hot Topics in Web Systems and Technologies, 2006. HOTWEB '06. 1st IEEE Workshop on*, 2006, pp. 1–11.
- [3] M. Castro, P. Druschel, A. M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "Splitstream: high-bandwidth multicast in cooperative environments," in *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*. New York, NY, USA: ACM, 2003, pp. 298–313.
- [4] J. J. D. Mol, D. H. J. Epema, and H. J. Sips, "The orchard algorithm: P2p multicasting without free-riding," in *P2P '06: Proceedings of the Sixth IEEE International Conference on Peer-to-Peer Computing*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 275–282.
- [5] V. Venkataraman, K. Yoshida, and P. Francis, "Chunkspread: Heterogeneous unstructured tree-based peer-to-peer multicast," in *Network Protocols, 2006. ICNP '06. Proceedings of the 2006 14th IEEE International Conference on*, 2006, pp. 2–11.
- [6] A. Datta, I. Stoica, and M. Franklin, "LagOver: Latency gradated overlays," in *ICDCS '07: Proceedings of the 27th International Conference on Distributed Computing Systems*. Washington, DC, USA: IEEE Computer Society, 2007.
- [7] M. Zhang, L. Sun, X. Xi, and S. Yang, "iGridMedia: Providing delay-guaranteed peer-to-peer live streaming service on internet," in *GLOBECOM, 2008*, pp. 1741–1745.
- [8] V. Venkataraman, P. Francis, and J. Calandrino, "Climber: An Incentive-based Resilient Peer-to-Peer," in *Proceedings of the Seventh International Workshop on Peer-to-Peer Systems (IPTPS 2008)*, 2008.
- [9] T. Ngan, D. Wallach, and P. Druschel, "Incentives-compatible peer-to-peer multicast," in *Proceeding of the 2nd Workshop on Economics of Peer-to-Peer Systems*, 2004.
- [10] BitTorrent, "http://bittorrent.com," 2004.
- [11] Gnutella, "http://www.gnutella.com," 2000.
- [12] eMule, "http://www.emule-project.net," 2002.
- [13] N. Andrade, M. Mowbray, A. Lima, G. Wagner, and M. Ripeanu, "Influences on cooperation in bittorrent communities," in *P2PECON '05: Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*. New York, NY, USA: ACM, 2005, pp. 111–115.
- [14] M. Feldman and J. Chuang, "Overcoming free-riding behavior in peer-to-peer systems," *SIGecom Exch.*, vol. 5, no. 4, pp. 41–50, 2005.
- [15] R. Axelrod and W. D. Hamilton, "The evolution of cooperation," *Science*, vol. 211, no. 4489, pp. 1390–1396, 1981.
- [16] A. R. Bhambe, C. Herley, and V. N. Padmanabhan, "Analyzing and improving a bittorrent networks performance mechanisms," in *Microsoft Research, MSR-TR-2005-03*, 2006.
- [17] G. Wu and T. cker Chiueh, "How efficient is BitTorrent?" vol. 6071, no. 1. SPIE, 2006, p. 607100.
- [18] S. Jun and M. Ahamad, "Incentives in bittorrent induce free riding," in *P2PECON '05: Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*. New York, NY, USA: ACM, 2005, pp. 116–121.
- [19] H. Park and M. van der Schaar, "Foresighted resource reciprocation strategies in p2p networks," in *GLOBECOM*. IEEE, 2008, pp. 2206–2210.
- [20] W. Myrvold and F. Ruskey, "Ranking and unranking permutations in linear time," *Inf. Process. Lett.*, vol. 79, no. 6, pp. 281–284, 2001.
- [21] M. Mares and M. Straka, "Linear-time ranking of permutations," in *ESA*, ser. Lecture Notes in Computer Science, L. Arge, M. Hoffmann, and E. Welzl, Eds., vol. 4698. Springer, 2007, pp. 187–193.
- [22] B. Cohen, "Incentives build robustness in bittorrent," in *Proc. of Workshop on the Economics of p2p Systems*, June 2003.