

Reducing Cost and Contention of P2P Live Streaming through Locality and Piece Selection

Wahidah Md. Shah, Johnathan Ishmael and Nicholas J.P. Race
School of Computing and Communications
Lancaster University
Email: [wahidah,ishmael,race]@comp.lancs.ac.uk

Abstract—The use of locality within peer-to-peer (P2P) networks is ensuring the construction of overlay networks that are both economically viable for network operators and scalable. However, the underlying protocols on which traditional P2P overlays are based are rapidly having to evolve in order to better support more time sensitive, real-time video delivery systems. This shift places greater demand on locality mechanisms to ensure the correct balance between bandwidth savings and successful timely playback. In this paper, we investigate the impact of peer locality within live streaming P2P systems and consider the pertinent challenges when designing locality based algorithms to support efficient P2P live streaming services. Based on our findings we propose an algorithm for supporting locality and harmonised play points in a live streaming P2P system. We present our results and in-depth analysis of its operation through a series of simulations which measure bandwidth consumption at network egress points, failure rates and each peer's play point relative to the live stream.

I. INTRODUCTION

Peer to Peer (P2P) was traditionally used as a technology to distribute large files on a best-effort basis. More recently P2P has become a technology through which video, both live and on-demand can be streamed from a single location to millions [1]. The overarching goal of the technology is to reduce the cost of distribution and deliver at the fastest speed. The technology forms an overlay network, without consideration for geography, network topology or cost associated with the data transfer.

P2P has a commercial impact to network operators due to the cost associated with traffic which traverses network transit points. Such behaviour have an impact on the overall cost to the ISP and saturate the egress link. This requires the network provider to implement traffic shaping to restrict P2P traffic, or to buy additional capacity on such transits. Such restrictions can apply to large ISPs, corporate network infrastructures and small regional networks such as Wireless Mesh Networks many of which will have a higher internal capacity than they do externally (to the Internet).

Previous research efforts has limit connectivity to between hosts on the same geographical areas or autonomous system (AS). However such restrictions can potentially impact the viability of P2P as a distribution mechanism as discussed in [2]. Namely, as P2P transforms into a system to facilitate live streamed media the question is raised, how does live-streamed media impact the success of locality and can additional metrics be introduced to overcome any potential impacts? This paper

considers P2P live streaming and the impact of introducing locality in order to limit the bandwidth crossing networks and the potential successes of this while considering how to maintain performance and quality of delivery, through establishing a unified playback point for geographically local, or local networked peers.

The rest of this paper is organised as follows: Section 2 presents an overview of the problem in locality live P2P streaming and the related work. The design and development of a simulation for experiment presented in Section 3. Finally in section V we present our conclusions and further work.

II. LOCALITY AND PIECE SELECTION

A. Problem Space

P2P networks are well known for their inefficient use of network resources due to the way in which they form an overlay topology on top of existing networking infrastructures. This uneconomical operation occurs during both content discovery and during the transfer of media itself when peers communicate with one another to transfer data. The inefficiencies lead to problems for network operators due the cost in routing the traffic and P2P networks/applications due to the sub-optimal nature of the network transfers. The cause of this issue stems from the way in which the P2P overlay network is formed, without consideration of the underlying network topology.

For an ISP this means P2P traffic may cross several inter-domain links typically resulting in increased costs, it also makes mitigation against such traffic (shaping) troublesome due to the dynamic, multi-peer nature of P2P. Approaches to relieving the problems caused by these problems usually resulted in a cat-and-mouse battle between ISPs and P2P protocol designers who typically change the P2P protocol to evade the restrictive mechanisms. More recently ISPs have been working to overcome these issues jointly, by introducing caches which locate a high capacity peer within an ISP to reduce the traffic crossing their network boundaries. P2P protocol designers have also been working with ISPs to help glean information on locality knowledge from the service provider. P2P protocols are also changing their function moving from a protocol for mass data downloads to delivering streamed media which restricts the viability of both caches, due to the time sensitive nature of the content and the speed at which overlays can be re-engineered to support locality.

The move to live streaming of media reduces the window over which content is both relevant and interesting to clients as most end-users wish to be experiencing the video content as close live as possible. This introduces several technical challenges which need to be addressed during the peer selection process such as selecting not only those with the highest capacity but also with a sensible playback point.

This raises questions related to understanding what interactions take place between streamed P2P media, locality and playback point and how can they be exploited to reduce delivery costs of streamed media for ISPs or Corporate Networks?

B. Related Work

There is an extensive set of work related to bandwidth reduction via locality in P2P networks these can loosely be categorised into the following research areas: **1)** Locality driven by ISP information, usually gained by an ISP by placing infrastructures within their network. **2)** Locality through dedicated *beacon* services running throughout the Internet (similar to 1, but without the Interaction with the ISPs). **3)** Through extracting information from other infrastructures (e.g. Content Delivery Networks (CDNs) DNS services). And finally **4)** Locality based on client side detection (e.g. Round Trip Time (RTT)).

In [3] Aggarwal *et al* present locality that provided by an 'Oracle' server that ranked peer based on metrics such as link delay and bandwidth estimations. Although there is significant improvement in overlay formation, they do not consider bandwidth savings. Similar are the iTrackers proposed by the P4P working group [4]. These provide trackers that operate at different network providers that provide locality information to peers.

The authors of [5] and [6] prevent cross-ISP traffic with biased neighbour selection policy, by modifying BitTorrent. Most of the selected neighbour peers are located within similar ISPs and permits a few peers outside ISP. They look at download times rather than streaming success. The work by Liu *et al* [7] differ from other work by considering the piece/item of content being transferred and how this may impact bandwidth consumption. Their findings show that enforcing locality with the BitTorrent choking mechanism does impact performance.

The use of Content Distribution Network technology is exploited by the authors of [8] who aim to reduce cross-network traffic cost by make use of DNS resolution technology employed by CDNs. They note that no great improvements are seen in download performance. CDNs are also manipulated by [9] who restrict peer into same city locations using CDNs technology.

While the work described above do not consider real-time media distribution, the work presented below do address the challenges related to locality and bandwidth efficiency within live streaming scenarios. We first start with the work presented in [10] who select peer that have better connectivity. Their results show a marked bandwidth reduction and upload capacity improvement, yet their results do not consider liveness as a factor. Liveness is considered a factor in [11], and their work

focuses around improvements to the start-up delay lag-from-live. A case study of PPLive is presented in [12] in which the authors discuss how the neighbour referral peer selection employed within PPLive has indirectly led to localisation, however no cost savings could be readily identified. In [13] the NAPA-WINE project discusses optimisation of video delivery and present a study on the level of network awareness in three P2P applications PPLive, SopCast and TVants their findings highlight little or no awareness of location in existing systems.

The previously presented work highlights the breadth and depth of research carried out in this space, while there has been significant work in reducing the cost of P2P transfers to network providers, typically this has been for non-sequential bulk transfers and not streamed media. Similarly work has been undertaken to improve locality in live-streamed media, yet this has typically focused on enabling closeness to live and not considering the bandwidth or cost implications. This paper seeks to consider both aspects; how to reduce the cost of P2P transfers through enabling locality whilst simultaneously focussing on the factors related to closeness-to-live for streamed media.

III. LOCALITY SIMULATION DESIGN

The impact of P2P streaming on network infrastructures was analysed through the use of a bespoke simulator, created using Python derived from the P2P-Next project¹. The decision to use a bespoke simulator was to enable a simple transition from demonstration environment to production (in-the-wild) code; without re-writing a significant proportion of the code base. It also enables validation of the simulator by testing it against the results from the code running in-the-wild.

The simulator is designed to test four modes of operation (1) Random - the traditional P2P model, in which no locality (2) Strict Locality - In which peers can only communicate with those local to it, or a seed if no other peers are available. (3) Variable Locality - In which a percentage of the peers are restricted to only communicating locally, while others can communicate with any peer (local or remote). And (4) Mixed locality, in which some peers honour the locality enforced upon them and others do not.

In the following sections we provide a description of the components which make up the simulator:

A. Tracker

The function of the tracker is to return a list of available peers and what pieces those peers hold back to an individual requesting peer. The number and type of peers returned from the tracker can vary depending upon the simulation configuration.

B. Peer

A peer represents an end-user, someone who is wanting to consume streamed media over the P2P network. When a peer first joins the network it first obtains a list of peers from the tracker. Next it chooses at what point to start streaming the file

¹<http://www.p2p-next.org/>

(known as the hook-in point). Then, the peer selects at random whom has the hook-in piece value and attempts to download it. A download attempt is only successful if there is sufficient AS link capacity (refer as contention ratio in the section IV) between the two peers. For example a peer contacting a remote peer on another network must have sufficient capacity on both the local AS link, remote AS link and at the peer. Otherwise, the download fails and re-download happened during the next iteration.

C. Locality

The locality model used for this simulation provides a straightforward approach to implement locality; each peer is designated to a specific AS. The tracker has four levels of localisation, first is random, which is no locality. Secondly, strict locality in which only peers in the same AS are returned or seed if no local peers are available. The third is random public and strict local (RPSL), in which strict locality is enforced by the peers on all but the first AS, with the first AS is keep random. This reflects a model in which not all the clients honour locality. The final locality is balanced locality, in which a percentage of neighbours are chosen to only talk locally, whereas other nodes are free to talk externally.

IV. EXPERIMENTATION AND ANALYSIS

The aim of our experimental analysis was to determine the impact of locality and play point on a live streaming P2P system. The overall aim is to reduce the bandwidth consumed at the Internet transit links between networks. We defined two simulation setups, mimicking two network environments. The first (*Scenario 1*), a corporate or community network in which internal connectivity between peers is high, secondly (*Scenario 2*) an ADSL style connection in which the upload capacity of each peer is poor. For both scenarios we define four different contention ratios (*AS20*, *AS50*, *AS220*, *AS440*), the value represents the maximum simultaneous pieces than can traverse the Internet transit at any one time. In each simulation there are 4 networks - 3 which represent end networks (e.g. ISPs) and 1 network which represents the public Internet. The AS capacity restrictions are only placed on the three ISP networks and not the link to the general Internet (which is unlimited).

Our first set of results presented in Figure 1 look at the typical P2P-Next code base when the capacity at each of the ASes are changed. With a heavily restricted AS capacity (**a**) the clients have insufficient bandwidth to download resulted into trends away play point from live. With higher capacity ASes (**b**) only some clients are able to stay close with live play point. The main reasons for peers failing to download a piece showing that peer capability is restricted by AS capacity in the smaller AS size and is limited by peer capability in the larger AS sizes.

In Figure 2 are the results with 50% locality enabled from both scenarios AS220. It show a marked improvement compared to the first experiment on the number of peers which are able to maintain a download. In graph (a), Figure 2

present some peer failure that futher away from live. On closer inspection they are mainly from AS3, suggesting locality has had a positive impact to two of the ASes, but not the third. Inspection to the reason of failure shows that majority of the failure rates are as a result of peer capability issues (which are worse with Scenario 2), interestingly the failures due to AS capability are much higher in Scenario 1, likely due to the greater capacity of the local peers, serving clients on the public Internet.


The scenario where every peer has locality enabled (not it is unlikely to happen in real network) does show the impact of locality should perfection be obtained. In Scenario 1 (which result are not shown) provides a near perfect play back for most clients which indicates that in a perfect environment, locality does permit playback of P2P media, even with highly contended AS links. The graphs in Figure 3 highlight that for Scenario 2 (a and b) they are unable to maintain a stream due to the limited capacity of the clients internal to the network.

Finally we look at a situation in which local peers have 100% locality, yet the public Internet peers have no locality enabled. Figure 4 highlights this behaviour is similar to that provided by 100% locality yet has a slightly higher failure rate due to the capacity of peers being taken up by those external to the network. The smaller AS capacities actually fair better as transfers fail due to AS capacity limits before reaching peers.


There is a complicated trade-off between the capability of an AS, the capability of its peers and the level of locality which is placed up a swarm. In situations in which AS capacity is limited then greater localisation is needed, but at a cost of greater capacity at the local peers. The graphs in Figure 6 show the bandwidth consumption at each of the ASes for selected simulation type. The initial spike at the start of each graph represents a new peer joining the swarm and the pre-buffer data. At 50% locality there is a significant saving (of up to 75%). Interestingly there is little difference between 50% locality and 100% locality with random public peers.

As part of our experimentation we wanted to consider the impact of selecting only local peers for a hook-in point thus ensuring local peers could service one-another and compare this to the standard P2P-Next hook-in mechanism. In Figure 5 we present the results comparing Scenario 1 and 2 with the modified hook-in point. The results from scenario 1 show that while there is a greater range in playback point than previous experiments all the experiments result in successful playback from all clients. Scenario 2 (with reduced client capacity) is worse in all cases (excluding random), than the same scenario without the new hook-in point.

When considering hook-in point, the behaviour can be considered as following; firstly peers joining the network will always pick plays points with a plentiful supply of content available locally (rather than globally), making it more likely that peers will communicate locally. The selection of the hook-in point will be directly related to failures within the local network, any peer struggling to advance its play point will lower the average for new clients. This reduces the demand on the peers with higher pieces, creating a fairer balance. In




(a) AS Capacity of 50




(b) AS Capacity of 220

Fig. 1. No Locality or Piece Selection: Each graph represents a single run of the simulator, configured to operate with a specific capacity limit for ASes 2-4 with no locality or piece selection enabled.




(a) Scenario 1 - AS Capacity of 220




(b) Scenario 2 - AS Capacity of 220

Fig. 2. 50% of the peers support locality, no piece selection.



(a) Scenario 2 - AS Capacity of 50



(b) Scenario 2 - AS Capacity of 220


Fig. 3. 100% of the peers support locality, no piece selection.

Scenario 1 which has ample internal connectivity this produces a near-ideal output for all scenarios, with clients are able to play back successfully albeit at the impact of being slightly further away from the live stream. In Scenario 2 all but the Random experiments fail due to the local clients being unable to support their own playback and are unwilling (due to the restrains placed upon them) to look externally. The clients therefore trend away from live and new clients then start further away from live. In the Random experiment, there is sufficient support from external clients to support the poor capabilities of the local peers, this coupled with the staggered playback is sufficient to allow the clients to play back.


V. CONCLUSIONS

This paper set out to expand on the promising work carried out in the space of locality within P2P systems. Specifically we were interested in considering how locality attempts to reduce the bandwidth which crosses network boundaries. Our focus was to consider real-time streaming and the implications this has for existing locality mechanisms. To achieve this we first considered existing systems, their purpose and operation. We then used a simulation environment, created based upon the P2P-Next code base to evaluate how locality impacted bandwidth consumption at network egress points, piece failure rates and a client's play point relative to live.

Our key conclusions and thus contributions are as follows:




(a) Scenario 2 - AS Capacity of 50




(b) Scenario 2 - AS Capacity of 220

Fig. 4. Random Public, Strict Local - The public Internet peers have no Locality, whereas the local peers enforce a strict locality.



(a) 50% Locality



(b) Random Public, 100% Locality Local

Fig. 6. Bandwidth Usage (Scenario 1 - AS Capacity of 220) - Each line represents the notional bandwidth carried over an AS, the bandwidth is measured in the form of pieces per simulation tick. Only selected graphs shown due to the limited space.


1) Locality mechanisms do provide significant bandwidth savings in live-streaming situations, however there is a strong trade-off between locality and the capability of peers (their available uplink capacity). A reduced client capacity results in significantly greater failures with locality enabled than without it. 2) Enabling a hook-in point which is aligned only to local neighbours, significantly improves the playback of P2P live streams with locality enabled, however capability of the local peers plays a crucial part in ensuring pieces are shared.

These findings provide a basis for further exploration of this subject, and raise a number of key questions not explored by this paper. Leaving aside the complexities of discovering what nodes are local to one another, how does a P2P system balance the many metrics required to enable locality successfully, who


and what determine what are suitable thresholds to enable successful playback, how is this system maintained (and at what computational cost).

REFERENCES


- [1] "PPLive, <http://www.pplive.com>."
- [2] F. Lehter, S. Oechsner, T. Hossfeld, Z. Despotovic, W. Kellerer, and M. Michel, "Can p2p-users benefit from locality-awareness?" in *Peer-to-Peer Computing (P2P)*, 2010 IEEE Tenth International Conference on, 2010, pp. 1–9.
- [3] V. Aggarwal, A. Feldmann, and C. Scheidele, "Can isps and p2p users cooperate for improved performance?" *ACM SIGCOMM Computer Communications Review*, vol. 37, no. 3, July 2007.
- [4] H. Xie, R. Y. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4p: provider portal for applications," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 351–362, 2008. [Online]. Available: <http://dx.doi.org/http://doi.acm.org/10.1145/1402946.1402999>
- [5] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving traffic locality in bittorrent via biased neighbor selection," in *ICDCS '06: Proceedings of the 26th IEEE International Conference on Distributed Computing Systems*. Washington, DC, USA: IEEE Computer Society, 2006, p. 66.
- [6] M. Alhaisoni, M. Ghanbari, and A. Liotta, "Localized multistreams for p2p streaming," in *International Journal of Digital Multimedia Broadcasting*, vol. 2010, 2010, p. 12.
- [7] B. Liu, Y. Cui, Y. Lu, and Y. Xue, "Locality-awareness in bittorrent-like p2p applications," *Multimedia, IEEE Transactions on*, vol. 11, no. 3, pp. 361–371, April 2009.
- [8] D. R. Choffnes and F. E. Bustamante, "Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 363–374, 2008.
- [9] C. Tian, X. Liu, H. Jiang, W. Liu, and Y. Wang, "Improving bittorrent traffic performance by exploiting geographic locality," in *Proceedings of the IEEE Global communications Conference*, 2008, pp. 1–5.
- [10] B. Biskupski, R. Cunningham, and R. Meier, "Improving throughput and node proximity of p2p live video streaming through overlay adaptation," *Multimedia, 2007. ISM 2007. Ninth IEEE International Symposium on*, pp. 245–253, December 2007.
- [11] E. Mykoniati, S. Spirou, R. Clegg, L. Latif, D. Griffin, and M. Rio, "Scalable peer-to-peer streaming for live entertainment content!" in *IEEE Communications Magazine*. IEEE, 2008.
- [12] Y. Liu, L. Guo, F. Li, and S. Chen, "A case study of traffic locality in internet p2p live streaming systems," in *ICDCS '09: Proceedings of the 29th IEEE International Conference on Distributed Computing Systems*. Washington, DC, USA: IEEE Computer Society, 2009, pp. 423–432.
- [13] D. Ciullo, M. Garcia, A. Horvath, E. Leonardi, M. Mellia, D. Rossi, M. Telek, and P. Veglia, "Network awareness of p2p live streaming applications: A measurement study," *Multimedia, IEEE Transactions on*, vol. 12, no. 1, pp. 54–63, 2010.




(a) Random - With PS




(b) Random - With PS




(c) 50% Locality - With PS




(d) 50% Locality - With PS




(e) Random @ AS1, Strict Local - With PS



(f) Random @ AS1, Strict Local - With PS



(g) Strict - With PS



(h) Strict - With PS

Fig. 5. Scenario 1 vs Scenario 2 - Piece Selection (Piece Picker)