On the Benefits of P2P Cache Capacity Allocation

Frank Lehrieder*, Valentino Pacifici‡, György Dán‡
*University of Würzburg, Institute of Computer Science, Würzburg, Germany
‡ KTH Royal Institute of Technology, School of Electrical Engineering, Stockholm, Sweden

***Abstract*—Peer-to-peer (P2P) systems are responsible for a large fraction of inter-ISP transit traffic in the Internet. Hence, many Internet service providers (ISPs) have deployed P2P caches to decrease their P2P related transit traffic. In our work we consider the problem of allocating the limited upload capacity of a P2P cache between a set of overlay swarms. The goal of cache capacity allocation is to increase the amount of transit traffic that can be saved using the cache. Our preliminary results are based on analytical models and simulations and show that cache capacity allocation is a promising means of improving the efficiency of P2P caches. We are currently validating our results via experiments on Planet-Lab.**

I. INTRODUCTION

Web caches have long been used by Internet service providers (ISPs) to decrease the amount of inter-ISP traffic and to decrease Web access latency. Since the available cache storage is typically insufficient to store all objects, much work has been done to develop cache eviction policies that would improve the cache hit rate, i.e., the probability that a Web request can be served from the cache [1].

The purpose of caches for peer-to-peer (P2P) systems is similar to that of Web caches, i.e., to decrease the amount of inter-ISP P2P traffic. P2P caches are widely used and were investigated heavily in recent years. Most of the works on P2P caches considered limited cache storage and were focusing on improving the cache hit rate [2], [3]. These works assumed that the transit traffic saving is equal to the amount of data served from the cache, and concluded that the cache hit rate can be improved by caching parts of the contents.

Peers can, however, download data from other peers even if it is available at the cache. As shown in [4], it is not straightforward to assess the impact of caches on the amount of inter-ISP traffic: the influence of the cache on the swarm dynamics and the peer selection mechanisms need to be taken into account to assess the actual benefit. As a consequence, the same cache upload capacity allocated to different swarms can lead to different savings in terms of transit traffic depending on the swarm parameters. Since the savings in terms of transit traffic is a function of the swarm parameters, it is important to understand whether the amount of transit traffic can be decreased by actively allocating the cache upload capacity to the peers in the swarms that coexist in an ISP's network. Understanding the impact of cache capacity allocation on the amount of transit traffic allows the design of more efficient P2P caches, and is an important step towards decreasing the tension between ISPs and P2P content distribution systems.

II. PROBLEM FORMULATION

Motivated by these observations we consider a scenario with a set $\mathcal{I} = \{1, 2, \ldots, I\}$ of I ISPs and a set $\mathcal{S}$ of swarms. Each ISP $i \in \mathcal{I}$ operates a cache with upload capacity K_i. We denote the cache capacity allocation of ISP i by the vector $\kappa_i = (\kappa_{i,1}, \ldots, \kappa_{i,S})$, with $\sum_{s \in \mathcal{S}} \kappa_{i,s} \leq K_i$. It is known that the cache upload capacity $\kappa_{i,s}$ allocated to swarm s has an impact on the transit traffic of ISP i [4]. It is, however, not well understood how the cache capacity allocation κ_i between swarm affects the amount of transit traffic of ISP i. Our goal is to answer this question.

Since ISPs are independent entities, the first question we have to answer is whether given the swarm parameters and the total cache upload capacities K_i there exists an equilibrium cache capacity allocation profile $(\kappa_i)_{i \in \mathcal{I}}$ from which no ISP would have an interest to deviate. To answer this question we model the problem of cache capacity allocation as a non-cooperative strategic game. Based on the observation that the amount of transit traffic savings for a swarm $s \in \mathcal{S}$ is a concave function of the amount of cache capacity $\kappa_{i,s}$ allocated to the swarm [4], we can show that the game has a pure-strategy Nash equilibrium, which by definition is an equilibrium allocation. The existence of an equilibrium cache capacity allocation indicates that ISPs could potentially settle in a stable allocation if the swarm parameters do not change.

III. CACHE CAPACITY ALLOCATION POLICIES

In order to evaluate whether cache upload capacity allocation can lead to a decrease of the inter-ISP P2P traffic, we consider various policies that ISPs could implement to allocate the available cache upload capacity among the swarms.

The baseline cache capacity allocation policy we consider is when the cache capacity is not actively managed by the ISP. The cache serves the requests of the peers according to a first in first out service discipline. As the request rate can exceed the maximum service rate of the cache, the cache is equipped with a drop tail queue in which requests can be stored until they get served. The cache capacity allocated to a swarm s is approximately proportional to the demand that the leechers of swarm s in ISP i put on the cache. Consequently, we refer to this baseline policy as the *demand-driven (DD)* policy.

Intuitively, if cache capacity allocation is to lead to decreased transit traffic compared to the demand-driven policy, then there must be some swarm s whose transit traffic decreases faster as a function of the amount of cache capacity $\kappa_{i,s}$ than that of the other swarms. If there is such a swarm s then the cache capacity $\kappa_{i,s}$ allocated to it should be increased until the marginal gain of allocating more cache capacity to it

978-0-9836283-0-9

This paper was peer reviewed by subject matter experts for publication in the Proceedings of ITC 2011

equals to the marginal gain of the other swarms. Nevertheless, finding the optimal allocation this way is difficult in practice, because the marginal gain of allocating more cache capacity to a swarm is hard to measure.

To approximate the theoretically optimal policy we consider policies that use priorities to allocate the cache capacity between the swarms. The priority of a swarm s is given by the ratio of the number of leechers $x_{i,s}$ in ISP i to the total number of peers in s. The cache serves the requests of peers using non-preemptive priority scheduling. We consider two ways of assigning priorities to swarms based on the ratio of local leechers and total number of peers. In the first case, swarms with lowest ratio have highest priority; we refer to this scheme as *inverse-ratio priority* (IRP). In the second case, swarms with highest ratio have highest priority; we refer to this scheme as *ratio priority* (RP).

As a comparison for the two priority-based policies we consider policies that reserve a fraction of the cache capacity to individual swarms. Such policies were considered in [5]. We consider three capacity reservation schemes. The first scheme reserves the same amount of cache capacity to all swarms, referred to as *uniform capacity reservation*. The second and third schemes reserve capacity to swarm s proportional to the ratio of local and total number of peers in the swarm, analogous to the two priority-based policies.

IV. PERFORMANCE OF ALLOCATION POLICIES

Our performance evaluation is based on simulations and experiments with real BitTorrent clients over Planet-Lab. For the evaluation we consider a BitTorrent system consisting of $|\mathcal{S}| = 12$ swarms. Motivated by the results in [4] we consider a network topology consisting of two ISPs, called ISP 1 and ISP 2. Peers arrive to swarm s in ISP i according to a Poisson process with rate $\lambda_{i,s}$. After downloading the file, they remain in the swarm for an exponentially distributed seeding time.

We implemented the cache capacity allocation policies in the ProtoPeer simulator, which was also used in [4], [5]. We were first interested in understanding under which conditions allocation policies have an impact on the amount of transit traffic. We defined three scenarios with different distributions of the peers over the swarms and ISPs and measured the incoming transit traffic savings, i.e., the fraction of transit traffic that can be saved by installing a cache as compared to the case without a cache. In the first scenario the peers are distributed uniformly between swarms, i.e., $\lambda_s = \sum_i \lambda_{i,s} = \lambda / |\mathcal{S}|$, but the arrival intensity to ISP 2 is ten times higher than that to ISP 1 (scenario *uniform,1:10*). The incoming transit traffic savings for the *DD*, the *IRP*, and the *RP* policies are presented in Fig. 1. We observe that cache capacity allocation does not make a difference for this scenario. The same is true for the second set of columns, results obtained for a scenario when the arrival intensities λ_s follow a Zipf-distribution, i.e., when some swarms are more popular than others (scenario *zipf,1:10*). We only observe a difference in terms of the transit traffic savings in the third set of columns. The third set of columns corresponds to a scenario where the arrival intensities $\lambda_{1,s} = \lambda_{2,s}$ for two swarms, and for the other 10 swarms the intensities are 1:10 (scenario *uniform,1:1+1:10*). The results show that the IRP policy performs best; it outperforms the *DD* policy by almost 30 percent. The reason is that when the ratio of local leechers is low, peers are likely to download data from peers in other ISPs. Therefore, they should be prioritized by caches. These results indicate that upload capacity allocation is an efficient means of improving the efficiency of P2P caches.

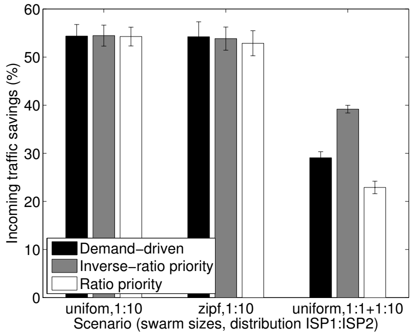

Fig. 1. Incoming transit traffic savings for 3 different scenarios.

Our simulations so far show that the *IRP* policy is superior compared to all other considered policies under a wide range of cache capacities, and work is ongoing on developing a near-optimal allocation policy to maximize the transit traffic savings due to caching. In order to show that the proposed cache capacity allocation policies can be implemented in real BitTorrent clients and caches and to evaluate the performance of the policies in a real environment we are currently working on experiments on Planet-Lab.

V. CONCLUSION

In this work we are investigating the impact of upload capacity allocation policies for P2P caches. Our main metric is the amount of transit traffic that can be saved due to caching. Our work is based on an analytical model of the system dynamics of BitTorrent-like systems and on simulations, and we are working on a Planet-Lab based validation of the proposed capacity allocation policies. Our preliminary results show that allocation policies can significantly increase the cache efficiency. Furthermore, the proposed priority-based policies are superior to policies based on capacity reservation.

REFERENCES

[1] N. Megiddo and D. Modha, "ARC: A self-tuning, low overhead replacement cache," in *Proc. of USENIX File & Storage Technologies Conference (FAST)*, March 2003.

[2] A. Wierzbicki, N. Leibowitz, M. Ripeanu, and R. Woźniak, "Cache replacement policies for P2P file sharing protocols," *Euro. Trans. on Telecomms.*, vol. 15, pp. 559–569, Nov. 2004.

[3] M. Hefeeda and O. Saleh, "Traffic modeling and proportional partial caching for peer-to-peer systems," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1447–1460, Dec. 2008.

[4] F. Lehrieder, G. Dán, T. Hoßfeld, S. Oechsner, and V. Singeorzan, "The impact of caching on BitTorrent-like peer-to-peer systems," in *Proc. IEEE Int'l Conf. Peer-to-Peer Computing (P2P)*, Aug. 2010.

[5] I. Papafili, G. D. Stamoulis, F. Lehrieder, B. Kleine, and S. Oechsner, "Cache capacity allocation to overlay swarms," in *Proc. Int'l Workshop on Self-Organizing Systems*, Feb. 2011.