

CLOSER: A Collaborative Locality-aware Overlay SERvice

Original

CLOSER: A Collaborative Locality-aware Overlay SERvice / PAPA MANZILLO, Marco; Ciminiera, Luigi; Marchetto, Guido; RISSO, FULVIO GIOVANNI OTTAVIO. - In: IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS. - ISSN 1045-9219. - STAMPA. - 23:6(2012), pp. 1030-1037. [10.1109/TPDS.2011.249]

Availability:

This version is available at: 11583/2439194 since:

Publisher:

IEEE

Published

DOI:10.1109/TPDS.2011.249

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

CLOSER: A Collaborative Locality-aware Overlay SERvice

Marco Papa Manzillo, Luigi Ciminiera, Guido Marchetto, Fulvio Rizzo, *Members, IEEE*

Author's version
Published in

IEEE Transactions on Parallel and Distributed Systems, vol. 23 n. 6, pp. 1030-1037

The final published version is accessible from here:

<http://dx.doi.org/10.1109/TPDS.2011.249>



Abstract—Current Peer-to-Peer (P2P) file sharing systems make use of a considerable percentage of Internet Service Providers (ISPs) bandwidth. This paper presents the Collaborative Locality-aware Overlay SERvice (CLOSER), an architecture that aims at lessening the usage of expensive international links by exploiting traffic locality (i.e., a resource is downloaded from the inside of the ISP whenever possible). The paper proves the effectiveness of CLOSER by analysis and simulation, also comparing this architecture with existing solutions for traffic locality in P2P systems. While savings on international links can be attractive for ISPs, it is necessary to offer some features that can be of interest for users to favor a wide adoption of the application. For this reason, CLOSER also introduces a privacy module that may arouse the users' interest and encourage them to switch to the new architecture.

Index Terms—P2P, file-sharing, traffic locality, privacy

1 INTRODUCTION

Motivation: Peer-to-Peer (P2P) file sharing systems have been experiencing a constantly increasing popularity during the last decade. This success is driving an evolution of these systems in terms of scalability, reliability, and decentralization. From Internet Service Providers (ISPs) point of view, file-sharing systems are both an opportunity and an issue: while these systems are a major driver for high-speed residential subscriptions, they force ISPs to increase their infrastructure bandwidth very often and, first of all, purchase more expensive transit services from Tier 1 carriers.

A promising approach to solve this problem consists in modifying one or more system components (e.g., the user application or the indexing system) in order to attempt directing requests to the closest peers that own

the requested resource (referred to as *resource providers* in the following). Examples of solutions adopting this method are presented in [?], [?], [?], [?], [?]. However, these solutions are suboptimal from a traffic locality perspective as in the selection of possible resource providers to contact for download they can consider only a subset of the available peers, thus potentially excluding some local providers. This is due to some design choices made in these solutions, which for scalability reasons cannot have access to the localization information of all the available resource providers (see Section 3.1 for details). Moreover, these systems do not give an adequate importance to the central role that users have in the evolutionary process of P2P systems. In fact, a significant percentage of the most widely used P2P applications has been developed and maintained by user communities, which need to be motivated to collaborate at the dissemination of novel systems and paradigms. At first sight, the locality-awareness seems to offer an intrinsic benefit for users that could stimulate their cooperation: since it reduces the average number of network hops crossed by download connections, it is statistically harder to traverse a bottleneck link and, consequently, the average download time should decrease. However, several publications [?], [?], [?], [?] demonstrate that this is not true in general and that in certain situations the download time may rather increase. Hence, we need different incentives which, similarly to the download time, are of interest for users. Without these incentives, the locality-aware techniques carry uneven advantages for ISPs and users, which may drastically limit their adoption in the existing P2P communities.

Contributions: Basing on these considerations, we developed CLOSER (Collaborative Locality-aware Overlay SERvice). CLOSER improves existing locality-aware solutions by offering the guarantee for downloads to be executed locally whenever is possible — i.e., when the

The authors are with the Dipartimento di Automatica e Informatica, Politecnico di Torino, 10129 Torino, Italy (email: name.surname@polito.it). ©2013 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

resource is present in the requester neighborhood — as it can discriminate among all possible resource providers when operating locality-aware selections. This is obtained with a negligible effort for the ISP and without affecting the scalability of the locality-awareness approach. The proposed solution is evaluated by both analysis and simulation, which also demonstrate the real importance of using the complete list of resource providers in the locality-awareness context. Furthermore, the feasibility and the simplicity of the approach are verified through the development of a real CLOSER-aware P2P application. CLOSER also introduces a novel mechanism to anonymize users' behavior in the network in order to stimulate their cooperation and hence favor the spread of the solution. This choice is motivated by the significant number of attempts to build *anonymous P2P systems* (e.g., [?], [?], [?], [?]) driven by open source, users supported, communities, which can be an indicator of users' vivid interest for privacy. These solutions are based on the utilization of proxy nodes as intermediaries during resource downloads, which guarantee anonymity, but to the detriment of the download speed [?]. The proposed privacy module overcomes these limitations by enabling direct downloads and it is shown not to violate the locality-awareness principle.

Outline: The paper is organized as follows. Section 2 presents the most prominent solutions applying the locality-awareness principles. The CLOSER architecture is described in Section 3, while Section 4 presents the privacy module that CLOSER includes to encourage users to change their P2P applications. Section 5 provides some analytical results concerning the effectiveness of CLOSER in lowering the inter-ISP link utilization, while Section 6 illustrates the simulation scenario and reports some simulation results showing the benefits stemming from the proposed solution. Finally, Section 7 concludes the paper.

2 RELATED WORK

Several possible solutions exist to provide traffic locality in P2P file-sharing systems. A promising approach consists in directing requests to the closest resource providers through the modification of some components of P2P systems, possibly in conjunction with the deployment of additional modules that slightly modify current P2P paradigms. Since also CLOSER belongs to such category, solutions adopting this method are briefly described in the following. Different approaches and their main drawbacks are instead presented in Appendix A, which can be found in the Supplementary File.

A first approach consists in modifying the behavior of current P2P applications, so that they can autonomously acquire their localization information and provide it to other users interested in the resources they share. In essence, a node acquires the list of resource providers from the indexing system. Then, it contacts the resource providers present in the list asking them for

their localization information and compares the obtained results with its own localization data. Closer resource providers are preferred to the distant ones. Examples of systems belonging to such category are Ono [?], a software extension of the Azureus BitTorrent client, and Kontiki [?], proposed in the context of P2P streaming. Each Ono instance determines its location by querying a Content Delivery Network (CDN) for a fake resource and collecting the mirror sites that the CDN chooses for it, according to the principle that users are redirected to a set of mirrors that are probably close to them (e.g., users always redirected to US mirrors are probably located in US). Kontiki implements a simpler localization methodology: starting from their IP addresses, Kontiki nodes obtain their AS Number (ASN) — assigned to ISPs by the Internet Assigned Numbers Authority (IANA) — from public databases.

A second approach consists in creating and exploiting a strong collaboration between users and ISPs, to be used in conjunction with some modifications to either the P2P application or the indexing system. In particular, each ISP deploys a special equipment providing the localization information to either the applications or the indexing system, depending on the specific solution. For example, [?] proposes to deploy a centralized equipment called *oracle* that users can query once they have acquired the list of resource providers from the indexing system. On the contrary, the P4P solution [?] proposes to deploy an *iTracker*, which is equivalent to the oracle facility but is directly contacted by the indexing system before sending the list of resource providers to a querying user. Thanks to the ISP collaboration, these techniques offer more precise localization information, with consequent improved performance in circumscribing traffic with respect to Ono and Kontiki. However, this results in an additional effort for the ISP, which has to deploy and maintain the equipment (oracle/iTracker) for ordering the list of possible service providers. It is also worth noticing how the presence of such equipment may allow malicious users to reconstruct the ISP topology — that generally is a confidential information — by forging ad hoc requests and analyzing the oracle/iTracker answers. On the other side, Ono and Kontiki do not require the ISP intervention and are less sensible to malicious peers aiming at reconstructing the ISP topology, but to the detriment of the localization precision.

More recently, these solutions have been used as a basis for additional work aiming at studying different aspects of the traffic locality problem. Considering a BitTorrent swarm, [?] and [?] investigated the impact of introducing the locality-aware principle not only in the neighbor selection, but also in the peer and piece selection procedures (i.e., the two operations which drive resource downloads in BitTorrent). Furthermore, [?] studied the adoption of BGP routing information as localization data used for ranking resource providers, while [?] explored the effects on the network when only a subset of resource lookups can be locality-aware. Also

the IETF expressed interest in the topic by forming an Application Layer Traffic Optimization (ALTO) working group [?] for standardizing a protocol for traffic locality in P2P systems. The IETF solution is based on an ALTO server which can be contacted to acquire the locality information, thus following the oracle/P4P approach.

3 CLOSER

3.1 Rationale

All solutions described in the previous section have a common operating principle: the locality information related to the resource providers is acquired (either by the applications or by the indexing system) whenever a user starts a lookup for a given resource, i.e., at *lookup time*. In Ono, Kontiki, and the oracle-based solution, users acquire a list of resource providers from the indexing system, and then collect locality information related to the listed providers. However, for scalability reasons, indexing systems generally do not supply nodes with an exhaustive list of resource providers, but randomly selects a subset of L resource providers among all the available ones — by default, $L = 50$ in BitTorrent. Let us denote this list as *sampled list*. Since every ISP includes a small percentage of the Internet population, it is unlikely that the sampled list includes a high number of resource providers located in the same ISP. Hence, the optimization process executed by these techniques may not be very effective. A similar problem is present in P4P, as the indexing system can send only a “sampled list” of the available resource providers to the iTracker, which hence performs a suboptimal ranking. Let us denote this issue as *sampled list problem*.

The rest of this section presents CLOSER, which avoids the sampled list issue as it ensures to consider all possible resource providers when discriminating among them. Furthermore, as it will be clearer in the following, CLOSER ensures the localization information adopted to be precise with a slight overhead for the ISP, which does not have to maintain any specific infrastructure.

3.2 CLOSER overview

CLOSER locality-awareness relies on two main principles: (i) the indexing system is made aware of the localization information of every resource provider, and (ii) this is done by enabling resource providers to communicate their localization information to the indexing system whenever they register a new resource, i.e., at *registration time*. During a lookup procedure, a requester gives its own localization data to the indexing system, which, thanks to these operating principles, can directly sort the resource provider list by increasing distance from the requester. In this way, even if the indexing system has to limit the resource list sent back to the requester to L entries for scalability reasons, the first L entries are the most interesting from the locality-awareness point of view. Hence, if the indexing system can guarantee to locate all the available resource providers (e.g., a

TABLE 1
Summary of modifications needed by locality-aware systems

System	ISP Support	Modified Application	P2P	Modified Indexing System
oracle	Required	Required		No
P4P	Required	No		Required
Ono	No	Required		No
Kontiki	No	Required		No
CLOSER	Optional	Required		Required

BitTorrent tracker or a DHT), it is possible to guarantee that if even a single resource provider is present within a given topological distance from the requester (e.g., in the same ISP or in the same country, depending on the adopted localization information), it will be sent to the requester with the correct associated distance. In essence, CLOSER enables a P2P system to discriminate among all possible resource providers without transferring the complete list. This is not possible with other solutions, which cannot use the complete list of providers for locality-awareness purposes as they should transfer the entire list over the network (generating the abovementioned scalability issues). Instead, our choice to move the localization data to the indexing system and their acquisition at registration time allows the locality-aware system to use the complete list in a simple and scalable way.

To make a P2P system CLOSER-aware, we need to modify both the indexing system — which has to be enabled to understand the localization information and sort the available resource providers according to this parameter — and the P2P application — which has to be able to interact with this new indexing system, referred in the following as *CLOSER indexing system*. Table 1 summarizes the modifications needed by current P2P systems to be compliant with the analyzed locality-aware techniques, including CLOSER. It is worth noticing that, although two separated columns are shown in the table for the modifications required by the P2P application and the indexing system, these two components coincide when the indexing system is distributed (e.g., a DHT such as Kademia [?] or in Gnutella [?]), as the indexing system is built and maintained by the application itself. This is the case for the majority of modern P2P systems (including BitTorrent), which tend to migrate to decentralized approaches. In CLOSER, the localization information has to be stored at the indexing system together with the resource itself. However, the small size of this information (a few bytes are sufficient to represent these data) makes this to result in a negligible cost if we consider the storage capabilities of modern computer architectures.

Additional details on CLOSER are provided in Appendix B and Appendix C, which can be found in the Supplementary File. In particular, Appendix B shows the resource registration and retrieve procedures, while

Appendix C describes the structure for the localization information we thought for CLOSER.

3.3 ISP support

In CLOSER, the resource providers themselves communicate their localization information to the indexing system during the registration procedure of new resources.

Similarly to what has been proposed for Kontiki [?] (see Section 2), resource providers may acquire their localization data autonomously (e.g., by querying public databases such as GeoIP [?]), without any intervention from the ISP¹. Although compliant with the CLOSER operating principles, this approach may reduce the accuracy of the localization information.

On the other hand, a proper ISP support can improve the system performance. In fact, if the ISP provides nodes with their localization information, this will result more accurate, thus allowing CLOSER to better achieve traffic locality.

To support CLOSER, ISPs do not have to deploy any infrastructure: they simply have to provide nodes with their localization data, which can be easily distributed through widely used systems — e.g., a web application. This is an advantage with respect to other techniques such as the oracle-based or P4P, which instead have to maintain specific servers. Furthermore, in CLOSER the ISP provides the localization information to each single resource provider and, thus, a malicious user should acquire localization data from every single user to reconstruct the ISP topology. This may be complicated as users' applications are not programmed to reply to direct queries concerning their topological information. This is another important difference with respect to the oracle/P4P scenario, where a malicious user can easily reconstruct the entire topology by simply interacting with the oracle/iTracker.

It is also interesting to remark that ISPs solely provide localization data; this produces some benefits: (i) future changes to P2P protocols do not require ISPs support and can be decided by the P2P application developers autonomously, solving the concerns highlighted by the P2P user community in [?]; (ii) users do not disclose information to ISPs or third parties, which the P2P user community highlighted as an issue in [?]; (iii) there is no legal concern for ISPs, because they do not participate actively either in the indexing system or in the resource exchange.

4 A USERS' PRIVACY MODULE FOR CLOSER

Section 3 focused on ISP needs, related to the circumscription of P2P traffic. Here we present CLOPS (CLOser Privacy Support), a privacy module for CLOSER that gives users an incentive to adopt CLOSER and hence to

1. These databases are useful for the higher level of the hierarchical localization information adopted in CLOSER. Lower hierarchical levels — e.g., the country and the town where the node is located — can be provided directly by users when starting the P2P application, since they usually know where they are located.

favor a wide spread of this new paradigm in the P2P community.

4.1 Rationale

The introduction of locality-awareness in P2P systems may clash with the indifference and the suspiciousness of users that, without proper incentives, are not motivated to adopt new P2P applications and paradigms. Furthermore, several publications [?], [?] demonstrate that locality-aware schemes may increase the download time, especially when peers are not uniformly distributed among the ISPs and their access bandwidths are heterogeneous. This leads to a *win-lose* situation for ISPs and users that further discourages users to adopt such systems. The oracle technique and P4P have an additional drawback from this perspective: users have to disclose information to ISPs, which are usually considered hostile [?], [?].

Hence, we need different incentives that, similarly to the download time, are of interest for users. Among the possible incentives, we select to focus on users' privacy, due to the effort that several user communities of software development are giving to the definition of *anonymous P2P systems* (e.g., ANts P2P [?], MUTE [?], OFF [?], Freenet [?]).

Using basic techniques, an eavesdropper that wants to compromise users' privacy can monitor their actions by (i) intercepting the control or data traffic generated by peers, (ii) acting as indexing system, by monitoring searches, shared resources, and downloads, and (iii) acting as a P2P user, by acquiring information during its apparently normal activity in the P2P overlay. The encryption features already deployed in modern P2P applications can easily prevent an eavesdropper to intercept P2P traffic and, consequently, this scenario is no longer interesting. Hence, we concentrate on the other privacy threats, which are of more interest in modern P2P systems.

The state of the art solutions concerning the users' behavior anonymity, also adopted in the abovementioned user-driven systems, is represented by [?] and [?], both proposed in 2002. These papers present two similar techniques based on the utilization of peers as proxy nodes, in conjunction with hard cryptography. A similar approach is also used in Tor [?], which has been specifically designed to anonymize TCP connections. These technologies make harder the connection tracing and, thus, hide who executes the requests, both when the eavesdropper controls the indexing system and when it acts as a normal P2P user. The penalty to pay when using these solutions is an increase of the download time [?], due to the utilization of possibly slow or overloaded intermediate proxy nodes to download resources whose size is generally large. This is perhaps the reason for which these techniques did not become widely popular. CLOPS, the users' privacy module we developed for CLOSER, avoids this issue by enabling direct downloads.

4.2 CLOPS overview

With respect to existing solutions, CLOPS achieves users’ privacy by following a totally different approach: P2P applications automatically select and download resources, even if those are not requested by the user. This offers users’ privacy because, observing a node behavior, it is hard to determine if resources were requested by the actual user or by an automatic download process. In particular, these automatic downloads can easily deceive an eavesdropper acting as a P2P user or indexing system as sharing or downloading a resource does not mean that the resource is shared or requested by the user.

In order to avoid penalizing actual traffic due to the consumption of precious access bandwidth of these additional downloads, it is necessary to introduce appropriate work-conserving scheduling algorithms that limit the CLOPS download rate. In particular, the downloading machine, which can discriminate among real and CLOPS downloads, gives higher priority to real downloads but ensures a minimum bandwidth guarantee to CLOPS downloads, selected as a small fraction of the total bandwidth available on the access link. In this way, CLOPS downloads can continue even when the access link is fully loaded, thus guaranteeing privacy, but do not penalize real downloads as in such a situation they consume a very small portion of the available bandwidth. The Class Based Queuing (CBQ) [?] algorithm can be used for this purpose as it is able to handle multiple classes at different priorities with minimum bandwidth guarantees. Hence, although possible alternatives exist, we propose to adopt the CBQ algorithm due to its proved effectiveness and wide adoption in many networking areas. Furthermore, several open-source CBQ implementations are available and can be seamlessly adapted to operate in the CLOSER context.

4.3 CLOPS content choice

Although from the privacy perspective CLOPS can select the resources to download in a random fashion, it could be worth investigating how the selection of such resources could influence the locality awareness of the system. In particular, techniques could be studied for favoring downloads of resources that may be of interest for the users of an ISP in a near future. In this way, users will be likely to download them from the inside of the ISP, thus improving traffic locality. However, our analytical and simulation results (presented in the following sections) shows how CLOSER by itself is able to keep local 98% of traffic, thus making any attempt to further investigate this aspect not very significant.

This considered, we just need to ensure that CLOPS downloads do not reduce the overall CLOSER performance. In fact, a completely random selection of resources clearly penalizes the locality properties of the system as resources may be downloaded from the outside of the ISP with high frequency. Hence, it is necessary to force CLOPS to download only a negligible percentage of resources placed outside the boundaries of the ISP.

TABLE 2
Model notation

Symbol	Meaning
N	# of resources in the P2P system
M	# of resources downloaded
$f(i)$	Probability that a user requests a resource of popularity rank i
$size(i)$	Size of resource of popularity rank i
P_{ISP_j}	Prob. user belongs to the ISP $_j$
P	# of users in the P2P system
Ω	Average # of shared resources per user
L	# of results obtainable by a real indexing system

Let p_{do} denote this percentage; a reasonable choice is $p_{do} = 0.1\% \div 1\%^2$. This policy can be applied thanks to the localization information offered by CLOSER, which enables CLOPS to discriminate between resources placed inside or outside the boundaries of the ISP.

In order to be able to select a resource to download, CLOPS modules have to be aware of the resources available in the P2P system. This is obtained by deploying a gossip protocol that spreads among nodes the information about the existence of resources. In essence, whenever a resource request arrives at the indexing system, this includes in its reply the ID of some resources randomly selected among the ones it knows. Analogously, whenever an interaction occurs between two nodes to start a download, those share the IDs of a subset of the resources they know. This enables nodes to learn existing resources and hence perform CLOPS downloads.

Appendix D, which can be found in the Supplementary File, describes a content encryption scheme that CLOPS adopts to avoid possible issues deriving from the presence of copyrighted or illegal material among the resources selected for automatic download. The appendix also details the algorithms adopted in a CLOPS-aware peer to perform both user-driven and automatic downloads.

5 A SIMPLE INTER-ISP TRAFFIC MODEL

Since inter-ISP links usually have the most significant associated cost, in this section we specifically focus on the performance of CLOSER in circumscribing P2P traffic within the ISP boundaries. In particular, we present a simple analytical model that shows how CLOSER outperforms not only the locality unaware systems (referred to as “LU” in the following), but also other locality-aware mechanisms (referred to as “ELA”) in achieving traffic reduction on inter-ISP links, thus also demonstrating the importance of the sampled list problem in the locality-awareness context.

Since each P2P protocol adopts different parallel download strategies (e.g., BitTorrent clients simultaneously download different file chunks according to specific piece and peer selection policies) and we would like to investigate a general case, we do not consider parallel

2. Notice that $p_{do} = 0$ affects users’ privacy as allows eavesdroppers to classify as real downloads the traffic exiting the ISP boundaries.

downloading in this model (analogously to the approach used in [?] for the oracle-based technique).

Due to space limitations, we present here the final outcomes of our analytical work. The complete analysis and some additional remarks are available in Appendix E, which can be found in the Supplementary File.

5.1 Traffic reduction on inter-ISP links

CLOSER/LU reduction. Given the notation described in Table 2, the percentage traffic reduction on inter-ISP links offered by CLOSER with respect to legacy systems can be obtained by

$$G_{C/L}\% = (1 - R_{C/L}) \cdot 100, \quad (1)$$

where

$$R_{C/L} = \frac{\sum_{i=1}^N s(i) \cdot f(i) \cdot (1 - P_{ISP_j})^{f(i) \cdot P \cdot \Omega} \cdot \text{size}(i)}{\sum_{i=1}^N s(i) \cdot f(i) \cdot (1 - P_{ISP_j}) \cdot \text{size}(i)}.$$

CLOSER/ELA reduction. Analogously to the previous case, we have

$$G_{C/E}\% = (1 - R_{C/E}) \cdot 100, \quad (2)$$

where

$$R_{C/E} = \frac{\sum_{i=1}^N s(i) \cdot f(i) \cdot (1 - P_{ISP_j})^{f(i) \cdot P \cdot \Omega} \cdot \text{size}(i)}{\sum_{i=1}^N s(i) \cdot f(i) \cdot (1 - P_{ISP_j})^{L_R(i)} \cdot \text{size}(i)}.$$

5.2 Traffic reduction evaluation

To quantify the real benefits of CLOSER in reducing the P2P traffic over inter-ISP links, we apply the above derived equations to a real-world case, adopting as a reference the network of Telecom Italia, a prominent Italian ISP. Data related to the Telecom Italia network that are of interest in this context are publicly available on the web. These are used to set the model parameters, as detailed in Appendix E, which can be found in the Supplementary File.

Under these assumptions, the percentage gain that CLOSER achieves with respect to both the traditional locality-unaware systems and the existing locality-aware solutions are reported in Table 3. We can observe how CLOSER guarantees about 98% gain with respect to LU systems and about 94.5% gain with respect to ELA mechanisms. These results demonstrate the effectiveness of CLOSER in reducing the utilization of inter-ISP links, thus making it an interesting solution for ISPs to limit their operating costs.

This result is achievable because a significant percentage of traffic is generated by popular resources that, by definition, are provided by a large number of resource providers. This effect is totally unexploited by locality unaware system, while the existing locality systems efficiency is compromised by the sampled list problem discussed in the previous sections.

TABLE 3
Traffic reduction on inter-ISP links.

Scenario	Value
Closer/Locality Unaware gain	97.9%
Closer/Existing Locality Aware gain	94.65%

6 SIMULATION AND EXPERIMENTAL RESULTS

Simulations have been run to both validate the above presented analytical model and further evaluate the proposed architecture. Some background on our simulation study and the setting methodology for the several parameters involved are presented in Appendix G, which can be found in the Supplementary File. All results are presented with 95% confidence interval.

In addition to this simulation study, we developed a CLOSER-aware application to verify the feasibility of our solution. Appendix H, which can be found in the Supplementary File, describes this software module and presents some results obtained on PlanetLab.

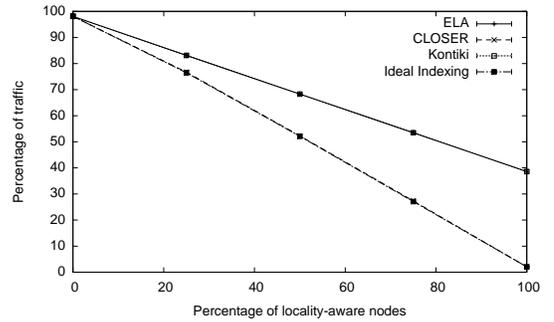


Fig. 1. Overall P2P traffic crossing the ISP borders

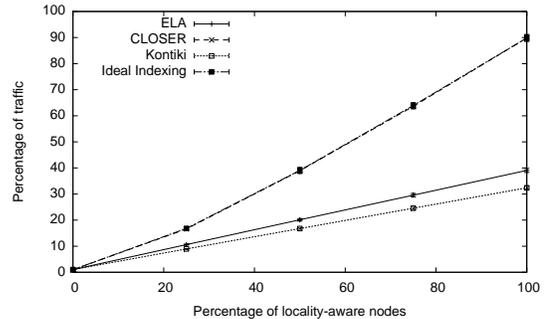


Fig. 2. Overall P2P traffic circumscribed to the requester's Area

6.1 Bandwidth usage

A first set of simulations aims at identifying how the available link capacity is utilized. We do not consider CLOPS downloads, whose effects on traffic locality will be presented later in this section.

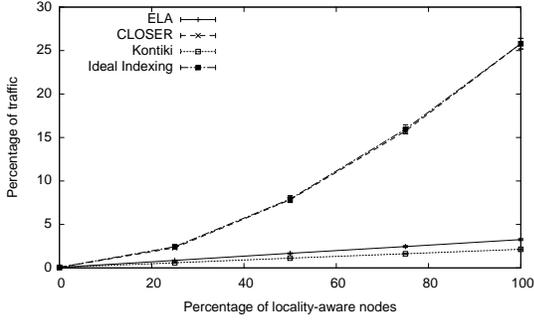


Fig. 3. Overall P2P traffic circumscribed to the requester's PoP

We evaluate the performance of: CLOSER, the newly described technique; LU, a generic legacy system without locality-awareness; Ideal Indexing, an ideal system that provides the whole list of content providers perfectly ordered according to the topological distance; ELA, the class of algorithms including Ono, the oracle, P4P, and ALTO; Kontiki, the simple mechanism used by Kontiki and described in Section 2. Kontiki and other ELA systems perform equal concerning the utilization of inter-ISP links, but they have to be handled separately in this simulation study as we also consider the circumscription of traffic within areas smaller than the entire ISP. In fact, Kontiki uses public IANA databases to acquire the localization information, which hence cannot be more specific than an AS number. Within the ISP boundaries, the resource provider selection of Kontiki is locality unaware, i.e., random.

Figure 1, Figure 2, and Figure 3 depict the link usage in different areas of the network as a function of the percentage of nodes in the P2P network adopting a locality-aware system. This is done to study the effects of a progressive adoption of locality-aware techniques.

Figure 1 reports on the usage of links with the tier 1 ISP. As expected, Kontiki and other ELA techniques have a similar behavior in this context (curves are overlapped in Figure 1). In fact, the effectiveness of both techniques is limited because of their ability of providing only a subset of the available resource providers (i.e., the sampled list problem). Both are outperformed by CLOSER, which mimics an ideal indexing system (again, curves are almost overlapped in the figure) thanks to its ability to offer the L closest content providers, perfectly ordered according to the topological distance. In fact, if a "local" resource provider exists, this will be included in the list and contacted by the querying user for downloading the file. If this peer is busy, the user will contact the next peer in the list, and so on until an available peer is found. Thereby, an ideal system providing a complete list of resource providers performs better than CLOSER only if more than L local resource providers exist and all of them are busy at the same time, which is an event unlikely to occur. Notice that CLOSER outperforms the ELA architectures despite the latter require

TABLE 4

Comparison of Simulation and Analytical results.

Scenario	Model	Simulation
$G_{E/L}$	60.74%	60.66%
$G_{C/E}$	94.65%	94.68%
$G_{C/L}$	97.90%	97.91 %

TABLE 5

Variation of inter-ISP link utilization due to CLOPS automatic downloads

p_{do}	Relative variation
0.10	-2.08 % \pm 1.18 %
0.25	-1.84 % \pm 1.19 %
0.50	-0.82 % \pm 1.34 %
0.75	-0.21 % \pm 1.18 %
1.00	0.81 % \pm 1.46 %

ISPs to deploy a powerful infrastructure composed by several servers. Table 4 compares the reduction of the inter-ISP link utilization obtained when the percentage of modified clients reaches 100% with the analytical results derived in Section 5.2, both confirming the effectiveness of CLOSER and validating our analytical model.

Figure 2 and Figure 3 show the amount of data that was circumscribed in an Area (the northern and southern Italy areas described above) and in a PoP, respectively. Also in these contexts CLOSER performs similar to an ideal indexing system, which confirms the effectiveness of the architecture also in handling the hierarchical localization information introduced in Appendix C.

6.2 CLOPS evaluation

To conclude our simulation study, we investigate the effects that CLOPS, the users' privacy module of CLOSER, has in the overall network performance. In particular, since CLOPS is based on automatic downloads, it is necessary to verify that this module does not affect the performance of CLOSER concerning the circumscription of traffic. Table 5 reports on the variation of inter-ISP link utilization due to the presence of CLOPS for different values of p_{do} (i.e., the percentage amount of resources that CLOPS downloads from the outside of the ISP). Although one could expect a performance degradation equal in percentage to the adopted p_{do} value, the table rather shows how we have a slight performance increase for small p_{do} values and a slight decrease when p_{do} grows. This is due to the presence, on average, of more copies of a resource within the boundaries of the ISP thanks to CLOPS downloads, which potentially lowers the utilization of inter-ISP links as reduces the probability for a user to download from the outside because internal providers are not available. However, since this event is unlikely to occur, CLOPS downloads results in a negligible increase of the system performance, especially when p_{do} grows. Aside these considerations, we can conclude that small values of p_{do} preserve user privacy and produce negligible effects on the utilization of inter-ISP links, which was our goal in this work.

The creation in the network of more copies of a given resource, due to CLOPS downloads, also explains the decrease of the average download time, although equal to 0.05%, we observed when CLOPS is used. As described in Section 4, a properly configured CBQ instance is introduced in the user machine to avoid penalizing real downloads due to CLOPS additional traffic (minimum bandwidth guaranteed to CLOPS download is fixed to 1% of the access bandwidth in these experiments). This considered, one probably expects an increase of the average download time, although slight thanks to the CBQ operation. Instead, the creation of more resource copies due to CLOPS increases the probability for a user to find a resource provider that is free and hence actually available to upload the requested resource. This lowers the average time that users' downloads have to wait in resource providers internal queues before being allowed to actually start, and consequently it lowers the average download time. This download time reduction (as said above, 0.05% decrease with respect to the system operating without CLOPS) is negligible. However, our real purpose was to avoid increasing this time, which is actually achieved in our system.

7 CONCLUSIONS

This paper presents the Collaborative Locality-aware Overlay SERvice (CLOSER), an architecture that aims at lessening the usage of expensive international links in P2P file-sharing systems. This is obtained by exploiting traffic locality (i.e., a resource is downloaded from the inside of the ISP whenever possible) and generates significant cost savings for ISPs. Analytical and simulation results show the effectiveness of CLOSER, also with respect to other proposed techniques for traffic locality in P2P systems. Unlike other approaches, CLOSER can discriminate among all possible resource providers, thus avoiding the *sampled list problem*. This is obtained without transferring the complete list over the network, thus also preserving the scalability of the system.

CLOSER also introduces a privacy module as an incentive for users to switch to the new architecture. Furthermore, a CLOSER-aware application has been developed and described in the paper.

Marco Papa is a Ph.D. student in Computer and System Engineering at the Department of Control and Computer Engineering of Politecnico di Torino (Technical University of Turin), Italy. He holds a B.S. Degree and M.S. Degree both in Computer Engineering. His research interests include quality of service, privacy and peer-to-peer technologies.

Luigi Ciminiera is professor of Computer Engineering at the Dipartimento di Automatica e Informatica of Politecnico di Torino, Italy. His research interests include grids and peer-to-peer networks, distributed software systems, and computer arithmetic. He is a coauthor of two international books and more than 100 contributions published in technical journals and conference proceedings. He is a member of the IEEE.

Guido Marchetto is a post-doctoral fellow at the Department of Control and Computer Engineering of Politecnico di Torino. He got his Ph.D. in Computer Engineering in April 2008 and his laurea degree in Telecommunications Engineering in April 2004, both from Politecnico di Torino. His research topics are peer-to-peer technologies, distributed services, and Voice over IP protocols. His interests include network protocols and network architectures.

Fulvio Rizzo is Assistant Professor at the Department of Control and Computer Engineering of Politecnico di Torino. He is author of several papers on quality of service, packet processing, network monitoring, and IPv6. Present research activity focuses on efficient packet processing, network analysis, network monitoring, and peer-to-peer overlays.