

Overlay Consolidation of ISP-Provided Preferences

Raul Landa, Eleni Mykoniati, David Griffin, Miguel Rio
Dept. of Electrical and Electronic Engineering
University College London
London, United Kingdom
rlanda, emykonia, dgriffin, mrio@ee.ucl.ac.uk

Nico Schwan, Ivica Rimac
Bell-Labs
Alcatel-Lucent
Stuttgart, Germany
nico.schwan, ivica.rimac@alcatel-lucent.com

Abstract—There is growing evidence that mutually beneficial outcomes can be achieved when content distribution overlays and their underlying ISPs collaborate through open interfaces.

However, most works in this area have focused on preference costs derived from symmetric network properties (e.g. RTT). The consideration of asymmetric preference costs is expected to benefit from the development of *consolidated topology construction* strategies that integrate the information provided by all participating ISPs and use it to produce an overlay topology with desirable global properties.

In this paper we propose a generic model for the multi-domain consolidation of ISP preferences expressed as cost-annotated lists of groups of topology-equivalent peers. Using this model, we propose two consolidated topology construction strategies: *Shared Cost*, designed to provide a tradeoff for preference cost asymmetries, and *Low Cost*, designed to reduce the overall preference cost that the overlay imposes on all its underlying ISPs. We evaluate these two models through extensive simulations over a wide range of ISP and PID size distributions, and we show that preference consolidation can provide ISPs with outcomes more aligned with their preferences than those provided by non-consolidated operation.

Keywords—overlay networks, ALTO cost maps, ISP preferences.

I. INTRODUCTION

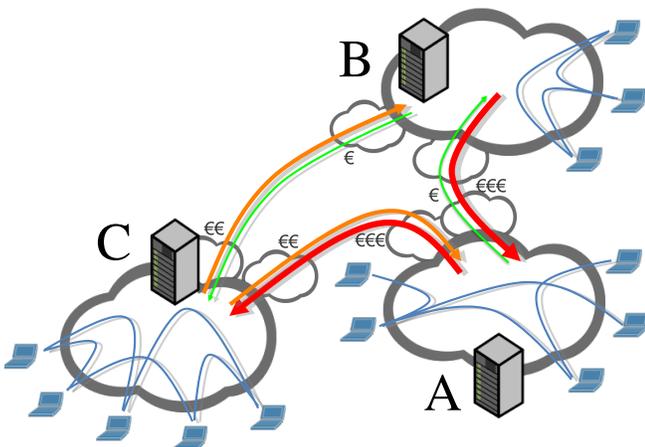


Fig. 1. Example: A possible source for asymmetric ISP preferences

Overlays collaborate with ISPs by driving their topology formation processes using information that each node obtains from its local ISP using open interfaces (i.e. [1]–[3]). This interaction usually involves ranked or annotated lists of network

regions called *PIDs*, which constitute clusters of topologically equivalent overlay nodes. By improving overlay construction via biased node selection, these collaboration techniques can be beneficial in reducing interdomain traffic and increasing overlay performance [3]–[8]. In this paper, we will denote Overlay-ISP collaboration interfaces as *OIC interfaces*.

Most studies to date have focused on using *network locality* as the basis for these annotated PID lists. However, OICs provide a good vehicle to implement other objectives such as reducing interdomain traffic costs or managing persistent traffic hotspots. These uses extend the study of Overlay/ISP collaboration into the realm of *asymmetric preferences*. Whereas locality-based OIC costs are symmetric, e.g. have the same properties in both directions, costs based on other network metrics may not have this property. For instance, ISPs contractually agree their interdomain costs with their service providers, leading to cost asymmetries. An overlay link between nodes in ISPs *A* and *B* may be *cheap* to *A* but *expensive* to *B*, or vice-versa. Therefore, a traffic allocation which is desirable to a given ISP may be undesirable for another one, and achieving an appropriate tradeoff therefore requires information from both ISPs. Taking interdomain cost as an example, Fig. 1 shows three multi-homed stub ISPs with arrows representative of the transit costs charged by their transport providers. Conventional operation drives *A* to bias overlay topology formation towards *B*, imposing upon it additional costs. Likewise, *C* could load balance equally between *A* and *B*, whereas biasing topology formation towards *B* and away from *A* would reduce the total cost.

We propose to address this problem through *consolidation of preferences*, a process whereby each OIC server provides a preference-annotated list of PIDs, which are collected by overlay peers and consolidated into a single preference-annotated list that represents an adequate tradeoff between the preferences of all ISPs involved. This list is then used by the overlay to drive its topology construction.

The first contribution of our paper is the definition of a generic model describing ISP preference consolidation, and is presented in §II. A further contribution can be found in §III, in which we present two preference consolidation strategies: *Shared Cost*, designed to provide a tradeoff for preference cost asymmetries, and *Low Cost*, designed to reduce the overall preference cost that the overlay imposes on all its underlying ISPs. We evaluate our consolidation strategies in

§IV, and show that preference consolidation can provide ISPs with outcomes more aligned with their preferences than those provided by non-consolidated operation. To achieve this, we present quantitative measures for preference alignment, and evaluate them through extensive simulations over a wide range of ISP and PID size distributions. Finally, we present the relevant related work in §V, and our conclusions in §VI.

II. PROBLEM DEFINITION

We consider an overlay with a presence in a set \mathcal{I} of N_I ISPs, so that $N_I = |\mathcal{I}|$. We focus in the construction of the *service topology*, that is, the set of overlay links that nodes use to exchange large amounts of traffic (rather than much smaller flows of signalling information). We call nodes adjacent in the service topology as *neighbours*, and consider the formation of this topology on the basis of ISP-provided information. We will assume that once a node has selected another node as a neighbour, this will trigger a unidirectional traffic stream at a standard average rate. Node selection will be our only consideration for topology construction; overlay-specific algorithms related to incentives or content availability (e.g. *tit-for-tat*) will not be considered.

We will assume that each overlay node will require k neighbours, and that each ISP will provide an OIC server that influences the topology formation choices of each node. To this end, we assume that each ISP i decomposes the entire set of overlay nodes into clusters of nodes which are considered essentially indistinguishable from the point of view of overlay topology formation; these will be called PIDs. We will assume that each ISP i defines both a set $\mathcal{P}_I(i)$ of *internal* PIDs that correspond its own customers, and a set of PIDs $\mathcal{P}_E(i)$ that comprises all other overlay nodes. Since these two are disjoint, each ISP can then define a set $\mathcal{P}(i) = \mathcal{P}_I(i) \cup \mathcal{P}_E(i)$ that represents all PIDs, both internal and external (of course, $\mathcal{P}_I(i) \subset \mathcal{P}(i)$ and $\mathcal{P}_E(i) \subset \mathcal{P}(i)$). We define the creation of an *overlay link* between PID l and PID m as the topology formation decision by a node in l selecting a node in m as its neighbour in the overlay. Then, we assume that each ISP i defines a *preference cost function* $c_i(l, m) : \mathcal{P}_I(i) \times \mathcal{P}(i) \rightarrow \mathbb{R}_+$ that assigns preference costs to overlay links. Note that we assume that an ISP i will only define preference costs for overlay links initiating in PIDs local to i (our model does not include transit ISPs, as in current OIC interfaces these do not announce preferences). Hence, $c_i(l, m)$ models how strongly does ISP i express a preference for nodes from an internal PID l to select nodes in PID m as neighbours in the overlay (m can be any PID, including internal ones). This preference cost function forms the basis for ISP-influenced overlay topology construction.

A. Preference Consolidation

In this paper we consider an multiple-ISP case in which each ISP i generates an inter-PID cost function $c_i(l, m)$. To simplify the problem, we will assume that there is a well-known decomposition of overlay nodes into PIDs, and that all PIDs are considered internal only for a single ISP. This means

that there is a universal \mathcal{P} so that $\mathcal{P}(i) = \mathcal{P}$ for all i , and that internal PID sets $\mathcal{P}_I(i)$ are disjoint subsets of \mathcal{P} . Hence, in $c_i(l, m)$, the ISP i is uniquely defined by the PID l , and we can directly use the shorthand $c(l, m)$ where l implicitly defines i through a mapping $i = \pi(l)$ where $\pi(l) : \mathcal{P} \rightarrow \mathcal{I}$ is a function that maps each PID l to its provider ISP i . Formally, we define the *generic cost function* $c(l, m) : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}_+$. This function maps any given pair of overlay PIDs l and m to a positive value representing the preference that ISP $\pi(l)$ has for nodes in PID l to create overlay links with nodes in m . To simplify matters further and aid comparison and aggregation, we assume that all ISPs agree on a common set of cost semantics.

Let the space of all generic cost functions $c(l, m)$ be called \mathcal{C} . A *consolidation function* is a function $F : \mathcal{C} \rightarrow \mathcal{C}$ that can be used to generate a single preference cost function from the partial preference cost functions provided by participant ISPs. Formally, if we denote our consolidated preference cost function as $\kappa(l, m)$, it follows that $\kappa(l, m) = F(c(l, m))$, where F can be designed so that $\kappa(l, m)$ has specific properties.

Since our objective is to compare these preference functions quantitatively, we require a baseline that represents default behaviour and from which we can measure any potential improvements. For this, we will use a F equal to the trivial identity mapping, so that $\kappa(l, m) = c(l, m)$. We will call this the *default consolidation*. Note that this corresponds to the case in which every node queries their local OIC server and ignores information provided by OIC servers of other nodes.

B. Overlay Topology Construction

We now address the modelling of how a particular consolidated cost function impacts the overlay topology formation process. We assume that each node will query the overlay for a set of candidate neighbours, and then select randomly from this set (this is reminiscent of OIC interaction, with the functions of the OIC server being provided by the overlay itself). The overlay can bias topology construction by providing candidate nodes in given PIDs with non-uniform probabilities. From a modelling standpoint, we simulate topology formation by determining how many nodes (and from which PIDs) will select nodes from a given PID as neighbours.

We start by defining a *population function* $N(l) : \mathcal{P} \rightarrow \mathbb{N}$ that assigns to each PID l the number of overlay nodes residing in it. In addition, we define a *topology construction function* $p(l, m) : \mathcal{P} \times \mathcal{P} \rightarrow [0, 1]$ that denotes the proportion of nodes in PID l that select nodes in PID m as their neighbours. Since $p(l, m)$ can also be interpreted as the probability of a node in l choosing a node in m as a neighbour, we have that $\sum_{m \in \mathcal{P}} p(l, m) = 1$. We will relate $p(l, m)$ with $\kappa(l, m)$ through a *preference-topology function* G , so that $p(l, m) = G(\kappa(l, m))$. Hence, G will model the impact that a given set of consolidated preference costs $\kappa(l, m)$ will have on $p(l, m)$.

III. CONSOLIDATED TOPOLOGY CONSTRUCTION

In the following we propose and evaluate two *consolidated topology construction* strategies: *Shared Cost (SC)* and *Low Cost (LC)*. Both these strategies consist of a consolidation

function F and a preference-topology function G . In addition, to serve as a baseline for comparison, we also describe a third algorithm that corresponds to a non-consolidated usage of the generic preference costs $c(l, m)$ provided by each OIC server.

A. Shared Cost (SC)

The objective of this consolidation strategy is to find a trade-off for the cost asymmetries present in the generic preference cost function $c(l, m)$. To achieve this, we define F as function that makes $\kappa(l, m)$ symmetric by assigning it the average of $c(l, m)$ and $c(m, l)$ weighed by the populations of m and l :

$$\kappa(l, m) = \kappa(m, l) = \frac{N(l)c(l, m) + N(m)c(m, l)}{N(l) + N(m)}. \quad (1)$$

We define G by considering $p(l, m)$ as inversely proportional to $\kappa(l, m)$, so that

$$p(l, m) = \frac{\kappa(l, m)^{-1}}{\sum_{n \in \mathcal{P}} \kappa(l, n)^{-1}}. \quad (2)$$

This means that, given a consolidated set of preference costs $\kappa(l, m)$, the overlay will select overlay links (l, m) with low $\kappa(l, m)$ with a higher frequency than those with high $\kappa(l, m)$.

B. Low Cost (LC)

The objective of this consolidation strategy is to induce an overlay topology that reduces the overall preference cost that the overlay imposes on all its underlying ISPs. To achieve this, the overlay will only create overlay links between PIDs l and m if $\kappa(l, m)$ is among the q lowest preference cost alternatives for PID l . More formally, we only allow nodes in a given PID l to become neighbours of nodes belonging to PIDs in a subset Q_l of cardinality q . We define Q_l to include the first q PIDs in a list of \mathcal{P} sorted by the preference cost $c(l, m)$. Hence, the PIDs in Q_l will be the q PIDs that have the lowest cost from l , and we have that

$$\kappa(l, m) = \begin{cases} 1 & \text{if } l \in Q_l \\ \infty & \text{otherwise.} \end{cases} \quad (3)$$

We define G in the same way as for the SC case, so that (2) still holds. For the specific case of LC, then, we have that

$$p(l, m) = \begin{cases} \frac{1}{q} & \text{if } l \in Q_l \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

C. Default

The objective of this consolidation strategy is to model the default situation where each node queries its local OIC server and performs node selection on the basis of this information only. For this strategy, we define F as the identity function so that $\kappa(l, m) = c(l, m)$, and use G so that (2) holds. Hence, $p(l, m)$ is inversely proportional to $c(l, m)$.

IV. EVALUATION

We will focus our evaluation on the degree to which the overlay behaviour induced by a given consolidation strategy aligns with the preferences of ISPs as expressed by their costs $c(l, m)$. We will start by defining a set of measures that will allow us to assess the performance of our proposed preference consolidation strategies.

A. Evaluation Measures

The first measure that we define is the *full Inter-PID preference cost* $f(l, m)$, that models the preference cost imposed upon PID l by all overlay links between l and m . Formally,

$$f(l, m) = k (N(l)p(l, m) + N(m)p(m, l)) c(l, m), \quad (5)$$

where $l \in \mathcal{P}$ denotes a *local* PID that has a number of overlay links with PID $m \in \mathcal{P}$, $N(l)$ the population of l , k the number of neighbours of each overlay node, and $p(l, m)$ the probability of a node in l choosing a neighbour in m . (see §II).

To model the preference cost that PID l incurs from overlay links both originating and terminating on its internal nodes, we simply aggregate $f(l, m)$ over m . This allows us to define the *ISP Cost* $\phi(i)$, that models the preference cost that ISP i incurs from overlay links both originating and terminating on its internal PIDs. Formally, we have that

$$\phi(i) = \sum_{l \in \mathcal{P}_I(i)} \sum_{n \in \mathcal{P}} f(l, n). \quad (6)$$

The expressions that we have just defined for $f(l, m)$ and $\phi(i)$ will allow us to estimate the overall preference cost that our proposed consolidation strategies have on the ISPs providing connectivity services to the overlay. We now provide a measure that directly assess the degree to which the outcome of a consolidation algorithm matches the preferences of each given ISP as expressed in their advertised OIC costs. To achieve this, we define the *preference concordance* $\delta(l)$, a measure for the *cosine similarity* between the full Inter-PID cost $f(l, m)$ and the generic preference cost $c(l, m)$ for a given PID l . We define $\delta(l)$ as

$$\delta(l) = \frac{\sum_{n \in \mathcal{P}} f(l, n)c(l, n)}{\sqrt{\sum_{n \in \mathcal{P}} f(l, n)^2} \sqrt{\sum_{n \in \mathcal{P}} c(l, n)^2}}. \quad (7)$$

Conceptually, $\delta(l)$ measures the cosine of the angle subtended between $f(l, m)$ and $c(l, m)$ if l is fixed and both are treated as vectors. Hence, if for PID l the full Inter-PID costs $f(l, m)$ induced by the topology construction function $p(l, m)$ have a similar structure as the ISP preference costs $c(l, m)$, we have that $\delta(l) \approx 1$ and the concordance is *high*. Conversely, if the structures of $f(l, m)$ and $c(l, m)$ are dissimilar, the costs induced by $p(l, m)$ are very different from the $c(l, m)$ expressed by the ISP, $\delta(l) \ll 1$ and the concordance is *low*.

B. Evaluation Environment

We use a simulation-based approach to evaluate a series of scenarios. In order to directly compare the results between simulation runs, the number of ISPs, PIDs and nodes was kept constant for all runs. In particular, all our simulation runs consider an overlay with $N_p = 10^6$ nodes, $N_P = 10^3$ PIDs, $N_I = 10^2$ ISPs, and $k = 10$ outgoing overlay links per node. However, within these constraints, we did vary the internal structure of the overlay by varying both the number of PIDs per ISP and the number of nodes per PID. In both cases, due to its simplicity and flexibility we used a Zipf distribution to assign sizes to ISPs and PIDs by their rank. This means that

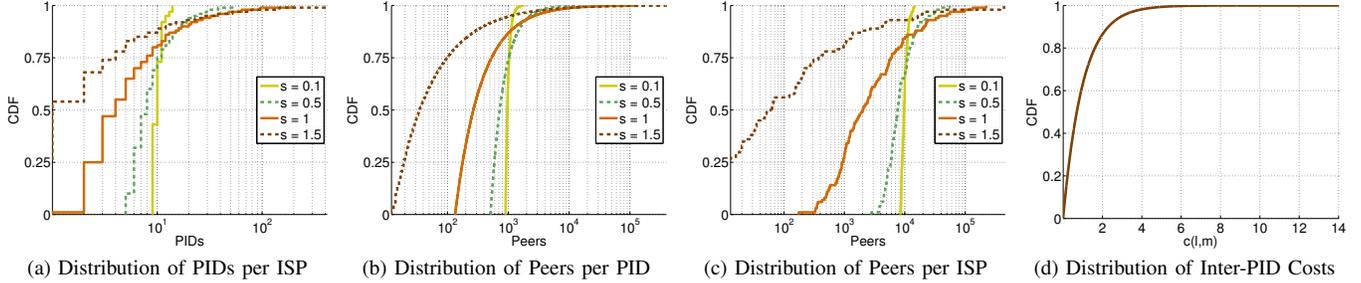


Fig. 2. Topology Construction Scenarios

the j -th biggest PID will have $N(j) = N_P z(j; s, N_P)$ nodes, where $z(j; s, N)$ is the Zipf density

$$z(j; s, N) = \frac{j^{-s}}{\sum_{n=1}^N n^{-s}}. \quad (8)$$

Similarly, the j -th biggest ISP will have $N_P z(j; s, N_I)$ PIDs. By changing the characteristic exponent s it is possible to explore different degrees of variability in the sizes of ISPs and PIDs. If s is small, $z(j; s, N)$ will be very similar for all ranks, and we have a situation where all ISPs have approximately the same number of PIDs, and where most PIDs have approximately the same number of nodes. Conversely, for larger s , $z(j; s, N)$ will decrease quickly with rank, leading to ISPs with widely varying numbers of PIDs and PIDs with widely varying numbers of nodes. We do not claim that this is a realistic model for a specific class of overlays. Rather, we focus our work on the analysis of the inherent properties of consolidation strategies, and use (8) as a controlled mechanism to generate a diverse set of overlay topologies for simulation.

In this paper we present results for a number of values of s in the interval $[.1, 1.5]$. Results for the values $s = \{.1, .5, 1, 1.5\}$ are shown in Fig. 2. The overlay structures induced by these values go from an almost homogeneous scenario in which there are approximately 10 PIDs per ISP, 10^3 nodes per PID, and 10^4 nodes per ISP (when $s = .1$) to a very heterogeneous scenario where these values range over 2 to 4 orders of magnitude across all ISPs and PIDs (for $s = 1.5$). Hence, the results presented are representative of a very wide variation on the ISP and PID structure of a large application-layer overlay. For the **Low Cost** consolidated topology construction strategy, we used $q = \frac{1}{10} N_P$.

With regards to the preference costs $c(l, m)$, we stress that they are not meant to be representative of actual interdomain billing costs. Instead, they are meant to convey the *relative* preferences of ISPs regarding which egress links to use. Since ISPs could potentially advertise their preference costs over widely different ranges, we assume that the overlay normalises them before processing so that $\sum_{n \in \mathcal{P}} c(l, n) = S_c$; this facilitates preference cost comparisons between PIDs and eliminates the problem of some PIDs reporting arbitrarily high costs. Given the lack of asymmetric preference cost datasets, we decided to generate $c(l, m)$ artificially by uniformly drawing values from the interval $[0, S_c]$ subject to the normalisation restriction; the resulting CDF is shown

in Fig. 2. In order for preference cost comparisons to be valid, the consolidated values $\kappa(l, m)$ were also normalised so that $\sum_{n \in \mathcal{P}} \kappa(l, n) = S_c$. For this paper, we used $S_c = 10^3$.

In order to quantify overlay-wide performance, we define two measures: the *total normalised overlay cost* Φ and the *total normalised preference concordance* Δ . These are designed to be overlay-wide versions of (6) and (7), and are defined as

$$\Phi = \frac{N_P}{N_P k S_c} \sum_{i \in \mathcal{I}} \phi(i), \quad \Delta = \frac{1}{N_P} \sum_{l \in \mathcal{P}} \delta(l).$$

Intuitively, Φ represents the overlay-wide average preference cost per overlay link once it has been normalised by $\frac{S_c}{N_P}$, the average $c(l, m)$ associated with any PID l ; Δ represents the average preference concordance over all PIDs in all ISPs.

C. Discussion of Results

The result of four sample simulation runs for $s = \{.1, .5, 1, 1.5\}$ are shown in Fig. 3. Each one of the curves in each graph correspond to one of the three consolidated topology construction strategies presented. In Fig. 3a, we see that for very homogeneous overlay structures ($s = .1$), the **Default** consolidation strategy produces consistently higher preference costs $\phi(i)$ for many ISPs when compared with both **SC** and **LC**, with **LC** producing the lowest preference cost. However, as shown in Fig. 3b, **LC** also produces overlays with very low concordance values $\delta(l)$, which may be unattractive for ISPs as they greatly deviate from their preferences $c(l, m)$. For these homogeneous overlays, **SC** not only achieves slightly lower preference cost, but also a greatly improved concordance of approximately .8. The **Default** consolidation strategy achieves a maximum concordance of around 0.4, and has a significant spread, indicating that the degree of alignment between $c(l, m)$ and $f(l, m)$ varies greatly between PIDs.

As overlay heterogeneity increases to $s=.5$ in Fig. 3c and $s=1$ in Fig. 3e, we see that the preference cost benefits provided by **LC** are reduced, since its $\phi(i)$ CDF extends further to the right. In addition, the CDFs for **SC** and **Default** increase their variability, including more ISPs with both smaller and larger $\phi(i)$. As a result of the increasing number of smaller ISPs, this manifests itself as a *shifting* of the CDFs towards the left. However, since a small number of much bigger ISPs are added as well, the expectation of the cost distributions exhibit limited change. Regarding the concordance $\delta(l)$, the CDF of **SC** is quickly shifted to the left first with $s = .5$ (Fig. 3d)

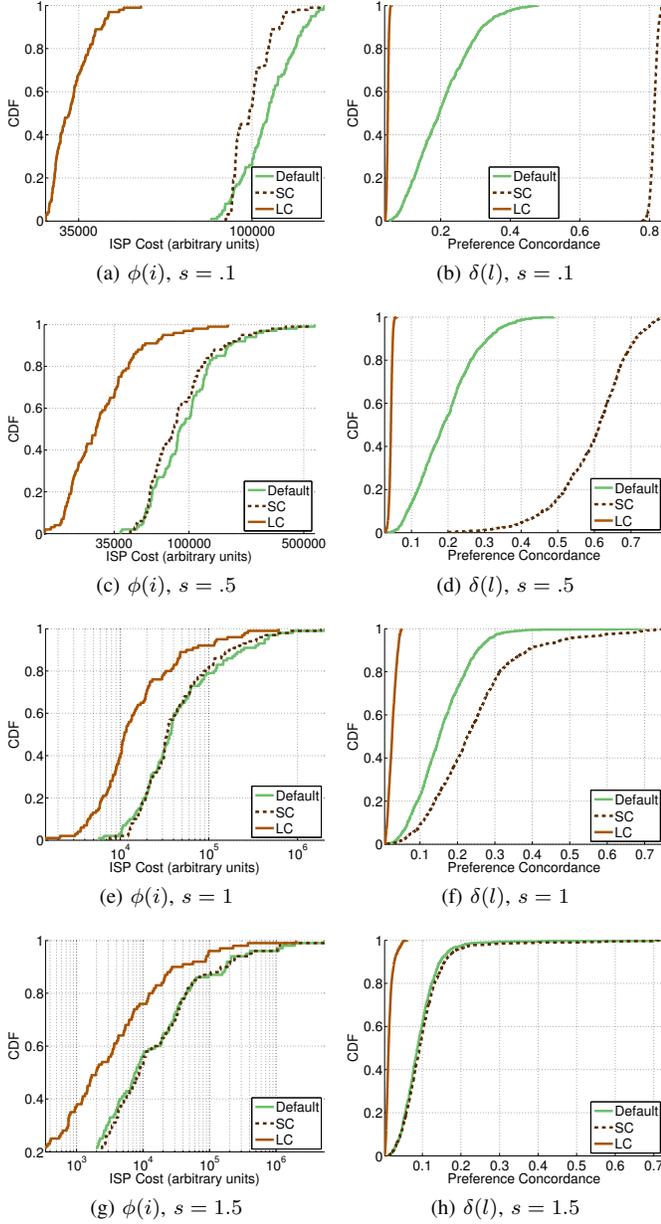


Fig. 3. Simulation results for $s = \{.1, .5, 1, 1.5\}$

and further still with $s=1$ (Fig. 3f), implying that as overlay heterogeneity increases the preference alignment of **SC** with $c(l, m)$ is quickly reduced for many PIDs. The concordance curve for **Default** is also shifted to the left, but at a much slower rate. The **LC** concordance experiences little change.

For the most heterogeneous case studied ($s = 1.5$), Fig. 3g shows little evidence for a change in the CDF of $\phi(i)$ for **LC**, **SC** or **Default**. However, the CDF of $\delta(l)$ continues shifting to the right, with **SC** and **Default** becoming indistinguishable.

In order to track changes in Φ and Δ more accurately, we performed ~ 36000 simulation runs for a sequence of values of $s \in [.1, 1.5]$. The results of these simulations are presented in Fig. 4. Each data point in Figs. 4a and 4b represents ~ 3000 simulation runs, with the error bars marking the 10-th and

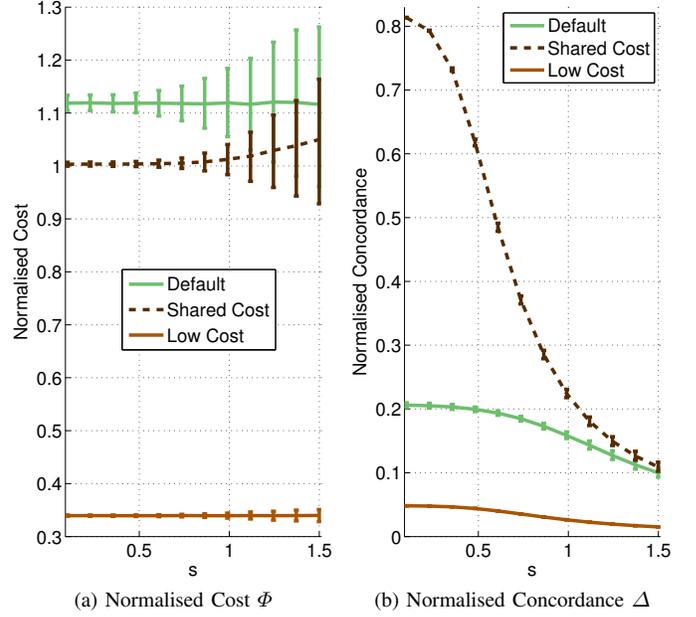


Fig. 4. Overlay-wide simulation results for $.1 \leq s \leq 1.5$

90-th percentiles of the Φ and Δ distributions respectively.

The general conclusions derived from the $\phi(i)$ and $\delta(l)$ CDFs in Fig. 3 are confirmed by this analysis. First, we note that as s increases, Φ for the **Default** and **LC** consolidation policies remains essentially unchanged in the mean, but exhibits increased variability. For **SC**, both Φ and the variability of the distribution increase slowly with s , with these increases becoming much more visible for $[1 \leq s \leq 1.5]$.

Finally, regarding Δ we see that for small values of s the **SC** policy has vastly higher concordances than those of either **LC** or **Default**. However, the expected Δ of **SC** drops very quickly as s increases, approaching that of **Default** for $s \geq 1.2$. This drop in concordance can also be seen for **Default** and **LC**, but these two cases it is much less pronounced.

The main conclusion that can be drawn from these results is that, for *all* values of s considered (i.e. ISP and PID size distributions), **SC** provides higher concordance Δ at a lower preference cost Φ when compared with **Default**. As shown in Fig. 3, these properties are also present in terms of the CDFs of both $\phi(i)$ and $\delta(l)$, with the $\delta(l)$ CDF of **SC** generally to the right of that of **Default**, and with the $\phi(i)$ CDF of **Default** generally to the right of that of **SC**. Hence, both when considering preference cost or concordance, **SC** is more compatible with the preferences of ISPs than **Default**. However, as shown in Fig. 4b this is critically dependent on the structural properties of the overlay, with **SC** performing much better for overlays with more homogeneous ISP and PID structures. A more detailed analysis of this issue, based on both overlay measurements and appropriate modelling based on the the Internet AS topology, is left for future work.

V. RELATED WORK

The expression of explicit interactions between overlays and ISPs is receiving increased attention by the research

community (see, for instance, [4], [9], [10]). One of the first works to focus in BitTorrent locality was [5], which relies on the ISP tagging each node as either *local* or *external*. This allows the overlay to set a soft limit on *external* nodes, biasing node selection and giving preference to intradomain connections. In [11], the authors present a system in which an *oracle* performs node ranking according to the preferences of the ISP. Although many possible metrics are presented as candidates (IGP metric distance, geographical information, expected delay, etc.) the authors focus in AS path length. In [4], this system is extended to take into account node upload bandwidth, and further improved upon by allowing the oracle to enrich DNS responses with ISP-provided information [12]. Another contribution in this topic is P4P [2], which became the basis for one of the main standardisation efforts in the area [1]. In its original presentation, [2] relied on the ISP aggregating nodes into groups called PIDs and then providing a set of end-to-end costs between the PIDs. These ISP-provided prices are used by the overlay to calculate its traffic matrix.

Regarding the reduction of ISP costs, multiple solutions have been proposed. Examples include ISP peering, IP multicast, content distribution networks and P2P localisation [6], [13], [14]. Another technique that has received increased attention from the research community is *traffic shaping*, that relies on reducing peak traffic volumes. This, in turn, reduces costs because traffic is usually billed on the basis of its 95-th percentile [15]. Particular approaches to reduce interdomain costs include simple rate-limiting/traffic shaping [16] or deferred transmission of delay-tolerant traffic [17]. Another proposal [18] proposes multiple ISPs to cooperate by jointly purchasing bulk transit bandwidth. This can be used to save costs due to the economies-of-scale effect of subadditive pricing as well as burstable billing: not all ISPs transit their peak traffic during the same period.

Finally, there has been some interest to approach the overlay-ISP interaction problem from the perspective of economics. As an example of this current of work, [19] presents a *market-centric*, architectural view. Additional examples of the analysis of cooperative outcomes between ISPs and overlays include the use of the Shapley value [20].

VI. CONCLUSIONS

In this paper we proposed a generic model for the consolidation of asymmetric ISP preferences expressed as cost-annotated lists of PIDs (e.g. topologically equivalent clusters of overlay nodes). Using this model, we proposed two consolidated topology construction strategies: **Shared Cost**, designed to provide a tradeoff for preference cost asymmetries, and **Low Cost**, designed to reduce the overall preference cost that the overlay imposes on all its underlying ISPs. We evaluate these two models through simulation, considering a wide range of ISP and PID size distributions. Our baseline for comparison was **Default**, a consolidation-free strategy in which each node biases its own topology formation algorithms using cost information that obtained only from its local ISP. We then compared our consolidated topology construction

strategies using two main tools: the *ISP Cost* $\phi(i)$, a measure for the preference cost that an ISP i incurs from overlay links both originating and terminating on its internal PIDs, and the *preference concordance* $\delta(l)$, a measure for the *cosine similarity* between the preference costs that an ISP experiences and its own preference costs as advertised through the OIC interface. Using these measures, we show that for asymmetric preference costs, consolidation strategies can provide ISPs with outcomes that are more aligned with their own preferences than those provided by non-consolidated Overlay/ISP collaboration.

This research has received funding from the Seventh Framework Programme (FP7/2007-2013) of the European Union, through the ENVISION project (grant agreement 248565).

REFERENCES

- [1] J. Peterson, E. Marocco, and V. Gurbani, "Application-Layer Traffic Optimization (ALTO) working group," 2009.
- [2] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider portal for applications," in *Proc. of ACM SIGCOMM*, 2008.
- [3] "ENVISION," 2009. [Online]. Available: <http://www.envision-project.org>
- [4] V. Aggarwal, O. Akonjang, and A. Feldmann, "Improving User and ISP Experience through ISP-aided P2P Locality," in *Proc. of the Global Internet Symposium*, 2008.
- [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 *Proc. of ICDCS '06*, 2006.
- [6] S. L. Blond, A. Legout, and W. Dabbousa, "Pushing bittorrent locality to the limit," *Comput. Netw.*, vol. 55, no. 3, 2010.
- [7] R. G. Clegg, D. Griffin, R. Landa, E. Mykoniati, and M. Rio, "The performance of locality-aware topologies for peer-to-peer live streaming (extended version)," *IET Software*, vol. 3, no. 6, pp. 470–479, 2009.
- [8] R. C. Rumín, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep diving into bittorrent locality," in *Proc. of INFOCOM*. IEEE, 2011, pp. 963–971.
- [9] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang, "Cooperative content distribution and traffic engineering," in *Proc. of NetEcon*, 2008.
- [10] N. Kamiyama, T. Mori, R. Kawahara, S. Harada, and H. Hasegawa, "ISP-Operated CDN," in *Proc. of the Global Internet Symposium*, 2009.
- [11] V. Aggarwal, A. Feldmann, and C. Scheidele, "Can ISPs and P2P users cooperate for improved performance?" *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 3, pp. 29–40, July 2007.
- [12] I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann, "Improving content delivery using provider-aided distance information," in *Proc. of IMC '10*, 2010, pp. 22–34.
- [13] M. Slot, P. Costa, G. Pierre, and V. Rai, "Zero-day reconciliation of bittorrent users with their ISPs," in *Proc. of Euro-Par 15*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 561–573.
- [14] D. R. Choffnes and F. E. Bustamante, "Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems," in *Proc. of SIGCOMM*, 2008, pp. 363–374.
- [15] X. Dimitropoulos, P. Hurlley, A. Kind, and M. P. Stoeklin, "On the 95-percentile billing method," in *Proc. of PAM*. Springer-Verlag, pp. 207–216.
- [16] M. Marcon, M. Dischinger, K. Gummadi, and A. Vahdat, "The local and global effects of traffic shaping in the internet," in *Proc. of IEEE COMSNETS*, Jan. 2011, pp. 1–10.
- [17] N. Laoutaris, G. Smaragdakis, P. Rodriguez, and R. Sundaram, "Delay tolerant bulk data transfers on the Internet," in *Proc. of ACM SIGMETRICS*, 2009.
- [18] R. Stanojevic, I. Castro, and S. Gorinsky, "CIPT: using tuangou to reduce IP transit costs," in *Proc. of ACM CONEXT*. ACM, 2011.
- [19] P. Faratin, "Economics of overlay networks: An industrial organization perspective on network economics," in *Proceedings of NetEcon*, 2007.
- [20] R. T. B. Ma, D. ming Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Internet economics: the use of Shapley value for ISP settlement," in *Proceedings of CoNEXT*, 2007, pp. 1–12.