On the Geography of X-Connects

Reza Motamedi
University of Oregon
motamedi@cs.uoregon.edu

Balakrishnan Chandrasekaran
Duke University
balac@cs.duke.edu

Bruce Maggs
Duke University / Akamai
bmm@cs.duke.edu

Reza Rejaie
University of Oregon
reza@cs.uoregon.edu

Walter Willinger
Niksun, Inc.
wwillinger@niksun.com

ABSTRACT

Companies like Equinix, CoreSite, and Telx manage and operate carrier-neutral colocation facilities (also called *colos*) where they provide, among other offerings, interconnection services. These facilities supply the infrastructure (*e.g.*, rack space, cabling, power, and physical security) necessary for network operators to colocate their routers for easy interconnection. This work focuses on identifying interconnections of the "cross connect" type, *i.e.*, dedicated point-to-point private peering links (which might be used to carry transit traffic or peer-to-peer traffic) that the network operators can buy from the colo providers so that their networks can exchange traffic within the confines of these facilities. In particular, our goal is to infer who is interconnecting with whom in which colos in which cities. Precisely locating the private peering links between two networks is a prerequisite for studying, for example, the root causes of the peering disputes between large content and eyeball providers in recent years. This paper introduces a multi-pronged approach for discovering these links. We illustrate the approach with case studies of colos in Seattle and Chicago. These studies demonstrate the promise as well as the challenges inherent in such a mapping effort.

Keywords

Internet Topology, Cross connect, Geography, Measurement

1. INTRODUCTION

To a first approximation, the decision about whether two autonomous systems (ASes) should peer boils down to economics and geography. Most importantly, the two parties must first determine whether peering makes sense financially and what form their relationship should take, *e.g.*, customer-provider or peer-peer. If a decision to peer is reached, the two parties must then select the locations at which the networks will interconnect. For example, they might start by looking for opportunities in which both networks already have a presence in the same building in the same city. While researchers have studied the logical construct known as the Internet's AS-level topology for over fifteen years, past efforts have focused almost exclusively on the existence and type of peerings and have largely ignored the physical realizations of these peerings. In this paper, we contribute to a better understanding of the Internet peering ecosystem by revealing the geographic nature of private peerings; that is, establishing which ASes buy cross-connects to interconnect with which other ASes in which colocation facilities in which cities.

Knowing where a given AS-level link is physically realized is paramount for studies that aim to shed light on the nature of Internet congestion (*e.g.*, [36] and references therein). In particular, to devise active probing campaigns that third-parties can perform for generating empirical evidence of "hot-running" peering links between two ASes, it is of critical importance to know exactly where these networks interconnect in a given city or region so that the probes target the correct cross connects. Operating hot-running peering links is a commonly-used practice in peering disputes between eyeball providers (*e.g.*, Comcast) and large content providers (*e.g.*, Netflix) whereby the former try to force the latter's hand. In these disputes, the eyeball provider may claim that it receives little or no additional benefit from pushing ever-increasing amounts of data from the content provider to its end users, and hopes that the poor performance experienced by the consumers of the content will obligate the content provider to pay for an upgrade of the interconnections (see [39, 58, 57, 53] for some of the recent peering disputes that have received attention in the popular press). How these disputes should be properly resolved is at the heart of the current debate about network neutrality.

Determining the geographic locations of cross connects is complementary to recent efforts that have focused on revealing the public peering links at the 350+ IXPs across the globe (see [12, 47] and references therein).[1] In this work, we concentrate on the US where, mainly for historical reasons, private peerings vastly outnumber public peerings [19, 20]. Due to this often-lamented scarcity of public peering opportunities (especially when compared to Europe), our focus on the geography of private peerings puts data center or colocation facility providers center stage.

---

[1] The discovery of private peerings utilized in IXPs remains an open problem, though [20].

Although largely neglected in prior studies of the AS-level Internet, in addition to providing colocation services to their customers, these companies are also in the business of selling interconnections that enable their customers to peer privately via a fiber cross connect or point-to-point circuit. As for-profit companies that are often also publicly traded, they are mandated to provide business-specific details in their SEC filings, and a cursory reading of some of their financial statements shows that Equinix dominates the cross connect market in the US with some 83K connections (as of 1Q 2015) [4], followed by companies such as Telx with some 48K and Coresite with about 15K [28, 33, 9, 3, 2]. To better understand the underlying economics, with an approximate monthly cost of $300 for a cross connect in any of the major cities in the US [1], in the case of Equinix, the revenues from selling cross connects in the US alone amount to about $250M per year [5].

Although these companies often publish lists of tenants in their colocation facilities, they generally do not indicate which tenants interconnect. Tenant lists can also be found in the widely used repository PeeringDB [45], but this database also does not indicate which clients within a facility interconnect. Furthermore, PeeringDB is incomplete. For example, for the colo facility we study in Chicago there are no entries in PeeringDB, and for the facility in Seattle the tenant list is inconsistent with the list published by the colo operator. The lack of public information about who connects with whom inside these facilities suggests that tenants wish to keep this information private. The current situation makes the discovery of cross connects at scale a challenging problem for Internet measurement, similar in spirit to the task of discovering the public peerings at IXPs worldwide [16, 50, 29].

The task of mapping the cross connects inside a commercial colocation facility, however, is inherently more challenging than mapping the public peerings between the different members of an IXP. At IXPs, member ASes use their IXP-assigned addresses to establish peerings with other member ASes. That these addresses are from prefixes that are assigned to and published by the various IXPs is at the heart of commonly-used heuristics for inferring the presence of a public peering link between two ASes that are members at an IXP based on information that is readily available from general-purpose traceroute probes [56, 16, 12]. The absence of any such comparable "hints" makes inferring the presence of an interconnection between two ASes that have PoPs in one and the same colocation facility extremely challenging and requires creative new solution methods.

In this paper, we report on a "baby step" towards our goal of mapping the cross connects in commercial colocation facilities at scale. In particular, we present case studies involving a colocation facility in Seattle, WA, and another in Chicago, IL, both of which are owned and operated by the Zayo Group. The Seattle facility is located in the Westin building, the premier carrier hotel in Seattle, with close proximity to numerous trans-Pacific cable landing stations in the Seattle area. Each of the two facilities has some 20 tenants and their names are published on Zayo's website. Using different measurement, analysis, visualization, and validation techniques, we show both the promise of our proposed approach when applied to mapping the cross connects inside these two colocation facility of the Zayo Group as well as the remaining challenges.

Our methodology consists of performing purposefully-designed and geographically-constrained measurements in the data plane (*e.g.*, targeted traceroute campaigns) and control plane (*e.g.*, selective looking-glass-based probes) and combining the obtained data with alternative non-measurement-based data sources. We then rely on domain knowledge, analysis, and visualization to deduce operationally and internally self-consistent connectivity structures that reflect Internet connectivity at the interface, point-of-presence (PoP), and autonomous system (AS) levels as seen and experienced from the perspective of a specific colo facility.

The remainder of this paper is organized as follows. Section 2 provides additional background information and discusses related work. In Section 3, we outline our methodology. Sections 4 and 5 explore how information gleaned from the control plane (*e.g.*, BGP data) can inform our search for the physical locations of peering links. Our data plane probing campaign is then described in Section 6. In Section 7 we explain our heuristics for mapping discovered peering links to specific facilities. The paper concludes in Section 8.

2. BACKGROUND AND RELATED WORK

Characterizing the topology of the Internet has been the focus of a large number of research efforts spanning almost two decades [44, 52, 18, 17]. While these studies have succeeded in unraveling some of the Internet's mysteries, solutions to certain problems (e.g., PoP-mapping) have remained elusive due to a variety of challenges related to performing Internet measurement studies. The challenges in topology mapping have been commonly attributed to the scale, heterogeneity and distributed nature of the Internet [30, 24, 41]; the lack of publicly accessible ground truth data also poses significant challenges, making it hard to evaluate tools and techniques. Various studies [42, 41, 15, 38] have pointed out the incompleteness in the measured or inferred Internet topologies and mention as a root cause the general lack of adequate Internet mapping tools and techniques.

Obtaining the Internet-wide interface-level topology hinges on the idea of performing large-scale traceroute measurement campaigns that are largely unconstrained with respect to the number and location of vantage points and targets. Internet topology discovery efforts have traditionally benefited from two types of measurement strategies. While the vast majority of measurements have been conducted in an opportunistic fashion, few have adopted a targeted approach to topology discovery. In the former approach, a large-scale measurement campaign is conducted to capture a presumably comprehensive view of the topology by relying on either a moderate ([18, 17, 37, 52, 37]) or large pool of vantage points [50] that serve as launching pads for performing many traceroute measurements. On the other hand, targeted measurements have been proposed to uncover smaller but more complete portions of the topology, by making informed decisions about the locations of the vantage points and targets in relation to the chosen portions. These studies are mostly motivated by the lesson learned from almost 20 years of large-scale Internet measurement studies which states "more is not always better" [49]. For instance, having Internet Exchange Points (IXP) in mind, Augustin *et al.* [16] used targeted traceroute probing to discover peering links at IXPs that are otherwise hard to detect. Durairajan *et al.* [25] use the triangle inequality property for their vantage point and target selection to increase the probability of observing PoPs that are missing from large-scale measurement campaigns.

The measurement methodology developed in this paper is motivated by such prior targeted measurement studies but is designed to tackle a much harder problem; that is, to shed light on the street-level geographic locations of the private peerings between pairs of ASes. In fact, while prefixes of an IXP or DNS names bearing a city name can be used to identify public peering at an IXP [16] or an unseen PoP [25] or interface [27], no such direct hints exist for geo-locating cross connects to a given colo facility in a given city. In the absence of such information, we rely on a diverse set of measurement-based indicators and non-measurement data that, when carefully combined, may provide sufficient information to place cross connects in the target colo or at least within the target city.

3. METHODOLOGY

Our goal is to infer the exact geographical location (*i.e.*, down to the building level) of physical realizations of logical AS-level links that we refer to as cross connects. The main intuition is that a majority of these cross connects are located in commercial colocation facilities. In this section, we present a systematic approach for discovering cross connects between ASes that are present in one form or another at a specific commercial colocation facility.

3.1 Connectivity Options in Colos

The international colocation marketplace is populated by a number of global companies (*e.g.*, Equinix, `equinix.com`) and numerous more region-specific businesses (*e.g.*, Cologix, `cologix.com` in North-America). To attract new customers and grow their business, these companies actively market and advertise their services and capabilities. In the process, they often publish detailed information about their colocation facilities on their web sites, including street addresses of the facilities, offered services, and customers/-partners and/or available carrier options. We use various public data sources (*e.g.*, PeeringDB [45], Packet Clearing House [43], data center map [23], open directory project [46]) and search engines to collect pertinent information for more than 1000 colo facilities in the US alone.

Given the highly competitive nature of the colocation business, different colo providers offer and support different connectivity options. The simplest option is for two tenants of a given colo to purchase a cross connect from the colo provider for directly connecting their routers. Here a "tenant" refers to a network that pays the colo provider for housing its equipment (*e.g.*, router) inside the colo facility, and this option places the purchased cross connect squarely inside the given colo facility. In addition to this simple case, there is also the option where a colo provider A (*e.g.*, Zayo [10]) partners with a more established colo provider B (*e.g.*, Equinix) to increase the connectivity options for its own tenants. To this end, A facilitates interconnecting its own tenants with tenants of B either via fiber risers (for partner colo providers in the same building) or by offering services such as Metro Interconnect (for partner colo providers with facilities in other locations across town). Such arrangements typically result in cross connects being purchased from B and established in B's facility, but A's reasoning for providing this option is that it prevents existing tenants from switching colo providers and has the potential to attract more business in the form of new tenants. To emphasize that B's tenant that is involved in such an arrangement with one of A's tenants is not a tenant of A, colo provider A refers to it as an "interconnection" or "carrier" option [10].

A similar situation occurs for large carriers that typically house their PoPs in their own buildings (*e.g.*, AT&T). If such a carrier wants to interconnect with some tenant in a given colo facility in a specific city, it often runs a local network service circuit from its own PoP in that city to its patch panel in the given colo facility and purchases a cross connect from that colo provider to establish connectivity between the patch panel and the tenant's router. This is yet another case where one of the two parties to a cross connect is not a tenant of the colo facility where the cross connect is located and utilized. To clarify the distinction between different type of tenants in a colo, in the rest of this paper, we refer to networks that are actual tenants of a given colo facility as *proper* tenants and to networks that are listed by a colo as "connectivity" or "carrier" options as *pseudo* tenants. We use the term "tenant" to collectively refer to both types of networks/ASes at a colo when the distinction is not critical.

Different colo facility companies have different approaches to publishing facility-specific tenant lists. For example, Equinix provides in general only minimal information in the form of total number of tenants and connectivity options and selectively-chosen customer testimonies. While it is possible to use PeeringDB-provided data to bootstrap the process of gathering somewhat more detailed information about specific Equinix facilities, this data is known to be rather incomplete and sometimes inaccurate or stale. At the other extreme, Cologix publishes the list of tenants for each of the markets where it operates colo facilities and also mentions additional connectivity options. The Zayo Group, on the other hand, typically lists the carrier options available in each of its facilities and also mentions additional connectivity options as a result of partnerships with various major colocation companies or the presence of IXPs in close proximity. In general, it is safe to assume that the Tier-1 ISPs that are listed as carrier options on these lists are not actual tenants but house their PoPs or routers elsewhere, either in their own facilities or in facilities owned and operated by some of the leading colocation companies. In contrast, many of the Tier-2 or regional ISPs are likely to be actual tenants, especially if they appear on the lists of the same colo provider in different markets across the US.

3.2 Our Approach in a Nutshell

The starting point for our approach is to mine the most promising data sources for obtaining colocation facility-specific information and extract a list of the tenants of a given colo facility; that is, networks that can either be identified as proper tenants or pseudo tenants (see earlier discussion).[2] Equipped with such a list, we then try to find cross connects within or in close proximity to the facility.

The basic idea is that since the colo in question houses the PoPs or routers of its tenant ASes, the latter have economic incentives to selectively connect to other tenant ASes (by purchasing cross connects from this colo provider) to support their business. At the same time, if the colo offers connectivity options to networks that, instead of being its own tenants, are tenants of some major close-by colo facility, we can expect to find cross connects that are in close proximity to the given colo. In this sense, our approach not only produces the cross connect matrix for the tenant ASes at the given facility but also generates information about cross connects that can be placed into facilities other than the target colo.

More precisely, our proposed approach is top-down and consists of the following four steps:

1. *Global and logical: AS-level connectivity between constituents.* We use readily available public BGP data sets to infer existing AS-level links between any pair of tenants in our target colo. Modulo known incompleteness and ambiguity issues, this analysis of global control plane information roughly defines the scope of the possible cross connects associated with this colo facility.
2. *Regional and logical: Geographically constrained AS-level connectivity between tenants.* We rely on purposefully-collected BGP information from vantage points that are in close proximity to the target colo to determine physical realizations of the logical AS-level links identified in Step 1 within the city or region where our target colo is located.
3. *Regional and physical: Geographically constrained physical connectivity between tenants.* Guided by the city/region-specific control plane information obtained in Step 2, we conduct purposefully-designed data plane measurements in the form of targeted traceroute campaigns to discover possible candidates for cross connects inside or in close proximity to our target colo.
4. *Colo-specific cross connects: Mapping the discovered cross connect candidates.* We leverage different features of our combined control and data plane measurements and rely on domain knowledge and various heuristics to determine whether the discovered cross connect candidates in Step 3 are are indeed located within the target colo or in close proximity.

In the following sections, we describe each of these four steps of our methodology in more detail. To facilitate the description of how our methodology works in practice, we use two specific colocation facilities as case studies for demonstrating the capabilities as well as limitations of our approach. Specifically, we select two facilities of zColo, a wholly owned subsidiary of the Zayo Group, a global provider of bandwidth infrastructure services and carrier-neutral colocation and interconnection: the zColo Seattle colocation facility (zSea) located at 2001 6th St. and the zColo Chicago colocation facility (zChi) located at 600 S. Federal St.

We consider these two facilities because of their moderate size and "challenging" environments. While zSea lists on-line some 20 networks as tenants (in the form of carrier options), the facility is located inside the Westin building, the premier carrier hotel in Seattle, and offers additional in-building connectivity options and also facilitates connectivity to both SIX (Seattle Internet Exchange) and the Pacific Northwest Gigapop (PNWGP). On the other hand, zChi provides a list of 16 tenants (also in the form of carrier options) and offers zColo's Metro Interconnect service to provide additional connectivity options to a large number of other colocation facility providers in the Chicago market.

Lastly, our methodology makes use of a number of more or less well-known tools and services. For example, for IP geo-location, instead if of using tools such as [35, 32] that are known to achieve poor accuracy for geo-locating IP address assigned to network infrastructure [51], we rely in this work on Akamai's professionally maintained IP-to-geo tool, called EdgeScape [13]. We also leverage information embedded in DNS names as network operators often encode geographic information in DNS names for their router interfaces [52]. Specifically, we use CAIDA's DDEC service [31] along with our own name-to-location mapping tool, called GINIE [14] to extract these hints from DNS names. Finally, we utilize Team Cymru's IP-to-ASN mapping service [54] to map individual IPs to different ASNs.

4. GLOBAL CONTROL PLANE VIEW

Control plane information (e.g., BGP table dumps) has

---

[2] While this task requires some detective work to identify the constituents' relevant ASNs, the methods are well-known and are not further discussed in this paper.

Table 1: *zSea* (zColo, Seattle) – Tenants of facility

AS Name	ASN	AS-Type	Target	Dist. (mi)	LG
AT&T	7018	T1	EdgeScape	1	-,-
BCE Nexxia	577	T2	EdgeScape	111	2,-
CenturyLink	209	Cable	Akamai	3	-,-
Charter Comm.	20115	Cable	EdgeScape	94	-,-
Cogent Comm.	174	T1	EdgeScape	1	-,-
Comcast	7922	Cable	Akamai	3	1,-
Earthlink	6983	T2	Prefix	1	-,-
Google	15169	Content	EdgeScape	1	-,-
GTT	3257	T1	EdgeScape	1	1,-
	4436	T1	EdgeScape	1	2,-
Hibernia Net.	5580	Regional	None	—	2,-
Level 3	3356	T1	EdgeScape	1	2,-
	3549	T1	EdgeScape	1	2,-
NTT Comm.	2914	T1	Akamai	3	-,-
PCCW Global	3491	T1	EdgeScape	1	-,1
Sprint	1239	T1	Akamai	3	2,-
TW Telecom	4323	T2	Akamai	3	2,-
Verizon Business	701	T1	EdgeScape	1	-,2
Windstream	6316	T2	EdgeScape	7	-,-
	7029	T2	Prefix	1	-,-
XO Comm.	2828	T2	EdgeScape	1	1,-
Zayo	6461	T2	EdgeScape	1	2,-
	5715	Regional	None	—	-,-

Table 2: *zChi* (zColo, Chicago) – Tenants of facility

AS Name	ASN	AS-Type	Target	Dist. (mi)	LG
Allstream	15290	Regional	EdgeScape	237	1,-
AT&T	7018	T1	EdgeScape	1	-,2
BCE Nexxia	577	Regional	Akamai	6	3,-
CenturyLink	209	Cable	Akamai	6	3,0
Cogent Comm.	174	T2	EdgeScape	1	2,-
Google	15169	Content	EdgeScape	1	-,-
Hibernia Net.	5580	Regional	EdgeScape	2	3,-
Level 3 Comm.	3356	T1	EdgeScape	1	3,-
	3549	T1	None	—	-,-
Lightower	46887	Regional	EdgeScape	1	-,-
TW Telecom	4323	Cable	EdgeScape	1	2,-
US Signal	26554	Regional	EdgeScape	2	-,-
Verizon Business	701	T1	EdgeScape	1	-,-
Windstream	6316	T2	None	—	-,-
	7029	T2	None	—	-,-
	1785	Regional	Akamai	6	1,-
XO Comm.	2828	T2	Akamai	6	3,-
Zayo	6461	T2	EdgeScape	2	3,-

been traditionally used for identifying a "global" view of the AS-level Internet where nodes represent individual ASes and links indicate the exchange of reachability information (via active BGP sessions) between nodes. However, as logical or virtual entities, AS-level links say nothing about the country, city, or building where this reachability information and, in turn, associated data traffic is being exchanged. In fact, an AS link can represent a single physical connection or hundreds of physical connections, established in one geographic location or in many different locations in a city, country, or across the globe. Nevertheless, by focusing on the list of tenants at a given target colo facility and utilizing readily available public control plane data, we can derive an estimate of the number of cross connects at that facility that is on the conservative side (i.e., not every inferred AS-link between tenant ASes is expected to yield a cross connect in this facility). At the same time, as a result of well-known incompleteness issues of the publicly available control plane data sets[42], it is possible to observe two tenants to be connected via a cross connect at a colo facility while the control plane data shows no AS-level link between the two ASes. However, such cases can be considered the exception and not the rule.

Tables 1 and 2 show the lists of tenants at the two target colos, zSea and zChi, and include details such as AS number (ASN), and AS-type (assigned manually based on public information available for each network); columns "Target", "Distance" and "LG" are explained later in Section 6. We combine these publicly available tenant lists with control plane data to determine which ASes are likely to have private peering relationships either at zSea or zChi, or both. Our control plane information is compiled from BGP data collected by RouteViews [11] and RIPE [8] and relies on inferred AS relationships from the Cyclops project [55]. This information was collected during January, 2015, and allows us to infer a matrix of pairwise relationships between the tenant ASes in zSea and zChi, respectively, for this time period. The matrix is omitted due to th elimited space but can be found in our related technical report [40]

Some care is required when generating such a matrix AS relationships, though. One the one hand, finding all ASNs associated with a tenant organization is critical to cross connect discovery since the organization can privately peer with others using one or more of its ASes. At the same time, naïvely including all ASNs associated with every tenant organization can quickly result in a huge matrix of candidate cross connects that needs to be explored and validated and can be expected to have many false positives. For instance, while there are two ASNs – 19092 and 6461 - associated with Zayo, AS19092 was not considered since it was observed only in logical AS links where the other endpoint was AS6461 and both these ASes belong to the same organization. In contrast, for Level 3, two ASNs – 3356 and 3549 – are included since both ASNs were observed to peer with multiple other tenant ASes at zSea and zChi.[3] As a result, Tables 1 and 2 show only those ASNs that were observed to peer with other tenants at the colo facilities.

Our discovery of potential cross connects at zSea and zChi distills data from numerous publicly available data sources and yields 253 possible links between the different ASNs used by the tenants in zSea and 153 in zChi. Exploiting the global control plane data view, we pruned 85 (33%) and 67 (44%) AS-links from the sets of possible cross connects in zSea and zChi, respectively. This leaves us with 168 and 86 potential cross connects in zSea and zChi, respectively. Note however that these are only rough estimates – not every inferred AS-link will be realized in these colos, and some physical connections may exist in these locations and give rise to AS-links that are missing from the global control plane view because of the incompleteness of the available

---

[3] Level 3 (AS3356) acquired Global Crossing (AS3549) back in 2011 but continues to use ASN 3549 to diversify its routing policies.

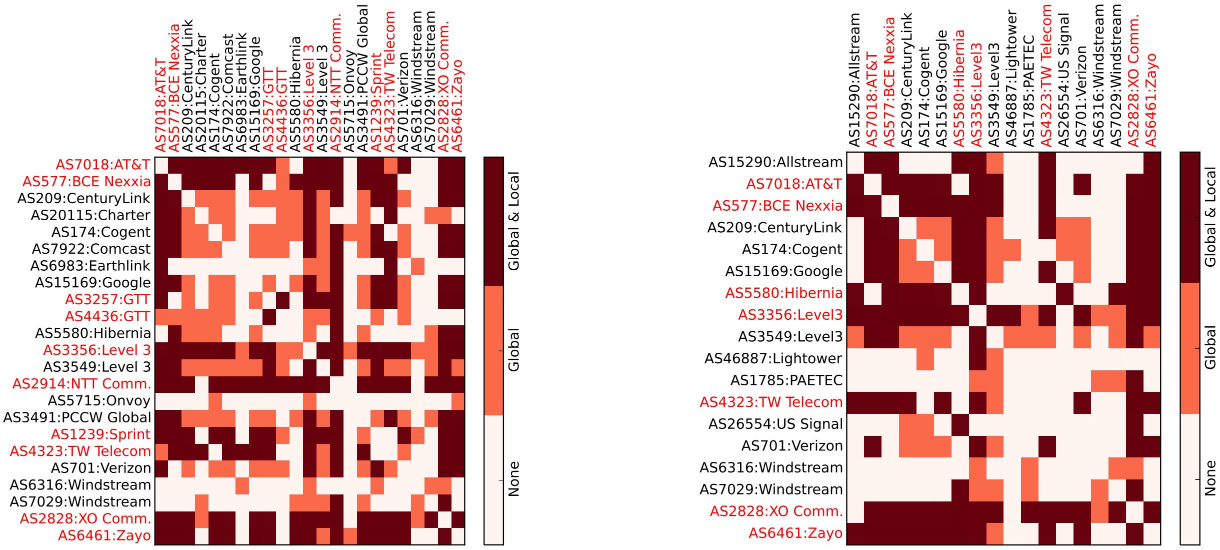

(a) *zSea* (b) *zChi*

Figure 1: Localized view of the control plane. The colors indicate the three cases of AS-links between tenants of the colo facility: (a) *None* – no relationship inferred, (b) *Global* – link found in global-view, but cannot be localized, and (c) *Global & Local* – link found in global-view, and can be localized using BGP records from a looking glass server. ASes in red contain one or more LGs in the same city as the target colo facility.

data.

5. LOCAL CONTROL PLANE VIEW

Based on a global view of AS-level connectivity enabled by the publicly available control plane data, we identified in the previous section the set of feasible AS-links that could represent cross connects used by the tenant ASes in our two colo facilities. Next we examine the local control plane view view of AS-level connectivity between tenant ASes. The reason for augmenting the global control plane view with its local counterpart is twofold. For one, while the global view might be missing some of the AS links between tenant ASes, a few or all of these links might be visible by the local view. Furthermore, being able to localize an AS-link to the city where a given colo facility is located increases the likelihood that the corresponding physical link representing a private peering between two tenants is in that city and possibly in the target facility.

Our localization efforts hinge on the existence and availability of BGP-enabled Looking Glass (LG) servers in the facility's city or close-by (both in geographic and network distance). Assuming AS_a provides access to LG_x, we can collect BGP records of the form $AS_a \rightarrow AS_b \rightarrow AS_c \rightarrow \cdots \rightarrow AS_n$ associated with some prefix p. While it would be unreasonable to expect that all AS-links in the AS-path provided by such BGP records can be localized, the chances that at least the first AS-link between AS_a and AS_b is in close proximity to LG_x are high. Thus, the use of LGs provides an opportunity to rely on control plane data to obtain a local (e.g., regional or city-level) BGP-view. Such a view can be very helpful for determining if an AS-link between two tenant ASes is physically realized in the same city or region where the colo facility is located. If the outcome is positive, it makes that AS-link a prime candidate for a cross connect within the given colo facility.

To localize the feasible AS-links identified in Section 4, we considered only those AS-links where at least one of the two involved ASes contains a BGP-enabled LG in the same city as the target colo. Particularly, in the case of localizing feasible AS-links to Seattle and Chicago, we identified a total of 17 BGP-enabled LGs (10 in Seattle and 7 in Chicago). These LGs can potentially cover a total of 129 AS-links associated with zSea and 66 with zChi from all AS links identified globally. Using the BGP records gathered from these LGs, we localized 99 (77% of the 129) AS links to Seattle, thus increasing the likelihood that these AS-links may be physically realized as cross connects inside zSea. Similarly, in Chicago, our heuristic localized 60 (91% of the 66) AS links. When grouping by inferred AS-link types, between $45\%-51\%$ of the localized AS-links are P2P, $38\%-45\%$ are provider-customer (C2P or P2C), and $\approx 10\%$ are unknown[4]. Note that while the fraction of localized provider-customer links is similar to that in the global view from Cyclops (see Section 4), we observe a higher percentage of P2P and lower

---

[4] Cyclops failed to identify the AS-relationship, but the AS-links still were localizable.

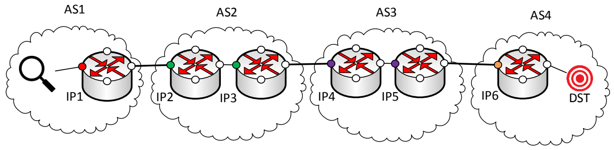

Figure 2: A traceroute and its corresponding AS-level path.

percentage of unknown AS-links.

Figure 1 shows the localized view of connectivity based on control plane information for tenants of zSea and zChi. The tenant ASes that provide BGP-enabled LGs are shown in red. In Figure 1, the relationship between AS_a and AS_b (*i.e.*, cell (a, b) of the matrix) is colored based on its visibility in the global and local views into three groups as follows: First, AS links that are revealed by both global and local views are colored in dark red. Second, AS links that are only visible in the global view are colored in light red. The local view may be missing an AS link for two reasons: i) when a BGP-enabled LG is available in any of the two ASes, these links may not be visible locally when local LGs provide the best (rather than all) AS path(es) to a prefix, ii) when no BGP-enabled LG was available in the related ASes to localize these links. Third, the AS pairs for which no AS link is present in the global and local views. Note that another possible case is AS links that are only visible in local view. While this case is generally feasible due to the incompleteness of the AS level maps, we have not observed any such link in our target colos. In summary, colored cells in Figure 1 show a union of all the AS links that were globally visible (as discussed in Section 4) and those that are locally visible. In the next section, we try to identify the physical cross connects corresponding to all these AS links.

6. LOCAL DATA PLANE VIEW

Our main objective in this section is to detect the physical realizations of the AS-links that result from our analysis in Section 5 and can (based on global and local control plane data) potentially be localized to the city or region where the target colo facility is located. In Section 7, we will describe our approach for determining if these physical connections indeed represent cross connects that can be placed inside the target colo or in close proximity to the target colo.

To explore how a logical AS-link that may be localized to the city or region where the target colo is located is realized as physical IP-level connections between pairs of tenants at the target colo, we conduct purposefully-designed data plane measurements in the form of traceroute probes that are intentionally biased to traverse cross connects that may exist in the target facility. A cross connect between two tenant ASes can be inferred from such traceroutes measurements by first mapping the individual per-hop IP interface addresses to their corresponding ASes (*e.g.*, using team Cymru's IP2ASN service [54]) and then examining the resulting AS-level path associated with such a traceroute for a link between a pair of tenant ASes. This is similar to the approach taken by prior work (*e.g.*, see [25, 50, 21]) and is illustrated in Figure 2 that shows the AS-level view of a (IP-level) traceroute and reveals various inter-AS links. Our ability to conduct such well-crafted traceroute measurements primarily depends on the the availability of measurement infrastructures with vantage points in "just the right" places. Note however that an inferred cross connect may not be at the target colo even if it is geo-located to the city of the target colo.

Given that the existing Internet measurement infrastructures all have more or less severe restrictions concerning vantage point selection, we leverage here three complementary strategies for performing traceroute measurement campaigns to partially offset these restrictions. We refer to the first strategy as **Scoped Probing**; that is, we conduct traceroutes between vantage points that are in close geo- or network-proximity of tenant ASes in the target city. The second strategy is called **BGP-Guided Probing**. In this case, we use BGP information to identify a pair of tenant ASes in the target facility where one party represents a regional AS and the other party serves as the regional AS's upstream provider. We then conduct traceroutes that traverse the resulting customer-provider link between this pair of ASes. The third strategy is called **Hosting Service Probing**, and its purpose is to reveal the cross connects associated with tenant ASes that offer web hosting services. This objective is achieved by identifying and probing web sites that are hosted by those ASes. All of these measurements where conducted during January of 2015. We next describe further details of our measurement campaigns that result from these three different strategies.

6.1 Scoped Probing

The basic idea behind this strategy is to conduct traceroute measurement between all pairs of tenant ASes (*i.e.*, criss-cross pattern in both directions) using vantage points that are as close (in geo- and network-distance) to the target colo as possible. This in turn increases the chances that the traceroute probes traverse cross connects that may exist at the target colo facility.

As the sources of our traceroute, we rely on public traceroute-enabled LGs provided by individual tenant ASes or their customer ASes (*e.g.*, based on the available AS maps from Cyclops [22]). For redundancy and to increase the yield, we use up to three LGs per tenant AS (or its customers). In addition to these LGs, we also use servers from the Akamai platform that are in the target cities. As targets, we only require a single destination IP in each of the tenant ASes that is geo-located as close to the target city as possible. Our options consist of (i) a list of more than 500 different Akamai servers across the US whose exact geographic locations are known, (ii) a list of 100M+ IP addresses in more than 15K different ASes [5] that been been geo-located using Akamais EdgeScape tool, and (iii) a hand-crafted list of

---

[5] `onrg.cs.uoregon.edu/Tau/`

Table 3: City-level inter-tenant connections identified by scoped traceroutes

	Seattle		*Chicago*	
AS link	**#links**	**#in-city**	**#links**	**#in-city**
Intended	302	120	178	67
Oppor.	234	59	77	20
Total Link	536	179	255	89

live IPs associated with prefixes that are advertized by the tenant ASes and are geolocated using the Maxmind geo-lite database [35]. Of all possible IPs, for each tenants AS we select the closet one to the target colo.

Due to rate limiting imposed by most LGs, care is required when launching traceroute probes from the LGs to the target IPs in each tenant AS to ensure that a campaign is completed within a reasonable time. In Tables 1 and 2, the last three columns provide the information about the best option for selecting source and destination for performing scoped traceroutes for our target colos in Seattle and Chicago, respectively. In particular, the columns labeled "target" and "distance" show the best available option for selecting the target IP and its distance from the colo in each tenant AS. The pair n, m in the column labeled "LG" indicates that we use n LGs in a tenant AS and m LGs in its customers. Note that all the selected LGs for an individual tenant AS have the same AS distance from the corresponding colo (*i.e.*, they are all located either in the tenant AS or in one of its customers).

We repeat each traceroute campaign 5 times during the second week of January 2015. We adopt a conservative approach and only consider the appearance of consecutive IP addresses in a traceroute probe that belong to two different tenant ASes as an indication of a direct inter-AS link (*e.g.*, cases where one or more IP-level hops appears between IP addresses of different tenants are not considered.[6] Note that a traceroute probe launched for the purpose of discovering a physical connection between two tenants may (unintentionally) discover an inter-AS link between some other pair of tenants.Obviously, such an "opportunistic" discovery is also useful and is utilized in our approach.

The combination of all selected LGs and target IPs at the tenant ASes results in 520 (Seattle) and 378 (Chicago) source-destination pairs for traceroute probing. However, only 422 and 306 probes successfully leave the origin ASes in zSea and zChi, respectively. Using Akamai vantage points we also can run 140 (Seattle) and 126 (Chicago) scoped probes, from which 131 and 89 leave the orgin AS succesfully.[7] Since we are primarily interested in inter-AS connections that are likely to be in the target city, we use EdgeScape and any geographic hints in DNS names to geolocate interfaces at both ends of each discovered inter-AS link. This information is then used to infer whether the discovered link is located in the target city or not. Table 3 presents the break down of all the identified inter-tenant links between tenants according to the location of the vantage point and target (labeled #links). If the link $AS_x \rightarrow AS_y$ is identified in a traceroute from AS_x to AS_y the link is considered as intended, otherwise it is counted as being opportunistically identified. Note that these sets are mutually exclusive. The table also reports this break down for the links that are geo-located to the target city (labeled in-city). As shown, the overall success rate of our probes in discovering physical connections (intended or opportunistic) between tenant ASes is high, but less than half of the discovered links appear to be in the target city. The heat maps that show the detailed view of the discovered links (all and the target in city) by our scoped probing technique are available in the related technical report [40]. Tables 4 and 5 summarize the main characteristics of the discovered links by this technique for zSea and zChi facilities.

6.2 BGP-Guided Probing

Our second probing strategy is to leverage the AS relationship between a pair of tenant ASes in the target colo when crafting traceroute probes for the purpose of discovering the corresponding physical link. The key idea is to identify any pair of tenant ASes that have customer-provider relationship and then launch traceroute probe toward an IP address that is served by the customer AS and is geo-located to the target city. The challenge here is that a relationship between an AS pair in the control plane could represent a number of different physical connections between the two ASes, possibly at different geographic locations. To increase the likelihood of identifying inter-AS links at (or in close proximity to) the target colo, we only consider pairs of tenant ASes where the customer AS has a regional coverage. The intuition is that such regional ASes establish physical connections to only a small number of provider/upstream ASes, typically at a single colo because managing different cross connects at multiple colos results in higher operational costs.

To perform BGP-guided probing, for each regional tenant AS_c, we first examine Routeviews data [11] to identify IP prefixes that are reachable through an AS path ending with $AS_p \rightarrow AS_c$, where AS_p is both a tenant AS and upstream provider of AS_c. We then scan the IP addresses in the identified prefixes and geo-locate them using Maxmind [35] to find a few live IP addresses that we map to our target city. These IP addresses serve as the targets for our BGP-guided traceroutes. Next, we launch traceroutes from a collection of vantage points towards these targets. Given the regional nature of the customer ASes and the chosen locations of the

---

[6] The presence of * hops in a traceroute sample may indicate traversal through hidden layer-3 infrastructure; therefore, it does not provide reliable evidence for a direct physical link between two tenant ASes.

[7] Further examinations revealed that the traceroutes that did not leave their origin ASes are associated with a few specific pairs of LG and destination IPs and were persistently blocked across all retries, and we discard these measurements.

Table 4: *zSea* - Summary of Data Plane View

		Interface		*IP Link*		*AS Link*	
Campaign	**#Probes**	**All**	**Exc.**	**All**	**Exc.**	**All**	**Exc.**
Scoped	3,307	154	86	179	140	49	25
BGP Guided	15,268	46	14	37	16	12	0
Web	3,648	103	47	107	78	26	6
Total	22,223	219		276		55	

Table 5: *zChi* - Summary of Data Plane View

		Interface		*IP Link*		*AS Link*	
Campaign	**#Probes**	**All**	**Exc.**	**All**	**Exc.**	**All**	**Exc.**
Scoped	2,618	101	69	89	74	33	20
BGP Guided	7,258	30	4	20	7	10	0
Web	777	26	8	20	10	8	0
Total	10,653	114		109		33	

traceroute targets in the target city, we expect that a significant fraction of the resulting traceroute probes traverse the physical connection between such (AS_p, AS_c) pairs in the target city if not at the target colo.

We identified 2 (5) regional ASes in zSea (zChi) to use for this technique. We initially launched our BGP-guided traceroute probes from LGs close to the locations of the Routeviews collectors that reported the corresponding AS-level paths towards the customer AS. However, since a majority of our traceroute-enabled LGs use UDP packets, most of these conventional traceroutes were blocked by middleboxes and did not reach their target in the customer AS. To address this issue, we ran TCP/ICMP traceroutes from three different large-scale platforms. In particular, we relied on Akamai Servers (*i.e.*, leveraged nearly 20 Akamai measurement servers in each of the two target cities), Planet Lab (*i.e.*, nearly 300 nodes [7]), and RIPE Atlas (*i.e.*, we used approximately 150 vantage points in a 500 miles radius of each target colo [48]). An overview of the number of discovered interfaces and links by this technique for each colo facility is reported in Tables 4 and 5. Although our BGP-guided probing discovers physical connections for the regional ASes, none of these intended connections is geo-located to the target city. Nevertheless, these measurement opportunistically identify 37 and 20 unique connections among tenants in Seattle and Chicago, respectively. The detailed statistics of the number of vantage points from the different platforms used to launch BGP-guided traceroutes along with the number of discovered links in each case are available in our technical report [40].

6.3 Web Probing

Our scoped and BGP-guided probing techniques are able to identify inter-AS links for tenants that are transit ASes or regional ASes. However, some tenant ASes are in the web hosting business, and for such ASes, the chances of finding their AS links as described above is low, unless we select the IP addresses of these web sites/services as targets for our traceroute probes. To identify such targets, we collect a large number of URLs that are associated with web sites hosted by the tenant web-hosting ASes in the target city. These URLs are collected by searching for local businesses websites (*e.g.*, from Yelp and Yellow Pages) and local schools and government branches websites available in various on-line directories. For each URL, we first identify the associated IP address, geo-locate it and map it to its ASN to select proper target addresses (*i.e.*, in the tenant AS and at the target city).

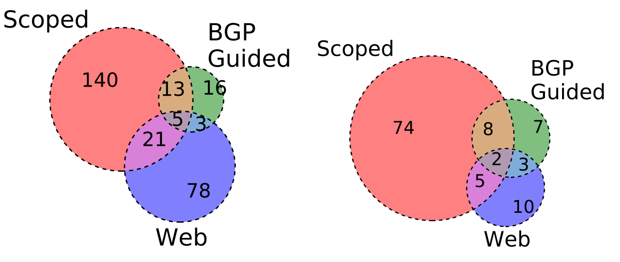

(a) Seattle (b) Chicago

Figure 3: Unique and overlapping physical connections discovered by different data plane measurement techniques

We identified 40 and 9 unique IP addresses that host a web page in a tenant AS of zSea and zChi, respectively. We then leverage the Akamai server platform (we use 7 (8) measurement servers, each in a different AS in Seattle (Chicago)) and RIPE Atlas (we use RIPE nodes [48] within a 500 mile radius of the target colos and pick the three closest nodes for each AS) to launch traceroutes to these IPs. This process resulted in the selection of 89 vantage points in Seattle and 83 vantage points in Chicago. Tables 4 and 5 summarize the number of all and in city links indentified by this technique for each colo facility. Detailed statistics are available in the related technical report [40].

6.4 Summary of Data Plane View of AS-Links

In this section, we summarize the result of all three data plane measurement techniques that we used to discover AS connections between tenant ASes. Tables 4 and 5 present the total number of probes, the number of discovered interfaces, the number of discovered IP links, and the corresponding AS links by each measurement technique (labeled as All). We also specified the number of the discovered entities that are exclusively discovered by each technique (labeled as Exc). Figure 3 provides further details on the number of links that were discovered by more than one technique. We observe that the scoped probing discovers the largest number of links and has the highest yield but other techniques have significant contributions as well.

In summary, all of our data plane measurement techniques find 276 and 109 interface-level AS connections that can be associated to physical connections in Seattle and Chicago, respectively. These interface-level connections correspond to 55 and 33 tenant AS pairs in zSea and zChi. All of these AS links are aligned with AS relations available in Cyclops and most of these AS relations (32 in Seattle and 21 in Chicago) are P2P. Figure 4 shows how the 276 and 109 interface-level AS connections in Seattle and Chicago are connecting various pairs of tenant ASes in the corresponding colo facility. We only consider these links for pinning in the next section.

7. MAPPING CROSS CONNECTS

Thus far the techniques presented for discovering cross connects between tenant ASes at best identify cross connects at city-level; the precise location of these cross connects at building-level within the city is yet unknown. In this section, we present a strategy to "pin" the interfaces at both ends of the discovered city-level cross connects to a specific colo facility in the city. It is worth noting that IP geolocation tools have much lower accuracy (on average at city-level or worse), and hence, can not be used for mapping interfaces to specific buildings.

The pinning technique examines any available evidence of the location of an interface either being inside or outside of the target colo facility, to accept (with some confidence) or reject the hypothesis that the interface in question is indeed at the target facility. In particular, we consider the following evidences to identify collocated interfaces:

- *Aliases:* We use iffinder alias resolution [34] on all the observed interfaces to determine interfaces that are likely to be associated with a single router, *i.e.*, each set of aliases represents a router.
- *PoP-tags:* Many organizations encode geographic location, PoP information, router type, and other details in the PTR records of their router IPs to aid in infrastructure management. A PoP-tag is a combination of three pieces of data: ASN, AS name, and PoP identifier. This information is extracted from the PTR records using our in-house PTR record parser in conjunction with DDEC [31]. For instance, the PoP-tag for `xe-10-0-0.edge1.Seattle3.Level3.net` is `AS3356:LEVEL3:SEATTLE3`. All interfaces with the same PoP-tag represent a PoP node.

In the context of cross connect discovery in zSea and zChi, iffinder identifies 13 IP addresses that form 5 groups of aliases in zSea and 8 IP addresses that form 4 groups in zChi. We also discovered 27 unique PoP-tags in Seattle and 16 in Chicago from PTR records of IPs; the PoP-tags in the median contain approximately 5 IPs[8].

The process of pinning inter-AS cross connects consists of two phases:

- *Identifying Anchors in a Colo*: we identify a set of physical interfaces (or node attributes, viz., PoP-tags) that we can confidently pin (or relate) to a colo facility in order to bootstrap the pinning process;
- *Inferring Colocation based on Association*: we use a set of *association rules* to relate the location of each unpinned interface n_i to one or more nodes n_j that are already pinned to the target colo.

To illustrate the basic idea of the pinning process, consider a scenario where there is sufficient evidence to conclude that interface n_x is located (*i.e.*, pinned) at a specific colo facility. If an interface n_y is associated with the interface n_x (*e.g.*, n_x is an alias of n_y or has a PoP tag similar to n_x), then the pinning procedure concludes that n_y has a high likelihood of being collocated with n_x. The subsequent sections present the two phases of the pinning procedure in detail.

7.1 Identifying Anchors in a Colo

Accurately identifying anchors in a target colo is crucial for our pinning procedure as they bootstrap the pinning process. This requires careful search and examination of available information for a colo that may directly reveal specific anchors (*e.g.*, an IP interface) or may offer an opportunity to capture anchor information through measurement. Such information include published data by network operators, scanning advertisements or marketing material from colo facility operators, identifying servers or vantage points at a facility, obtaining information from the operators of a colo facility or its tenant ASes. Clearly, the utilities of this information vary across different colo companies and facilities. In this section, we present two techniques for identifying anchors in the zColo facilities, zSea and zChi.

Using Online Information: For customers in Seattle and Chicago, Zayo Group (which owns the two facilities) offers a *Long Haul Dark Fiber* [6] enabling Zayo's customers in Chicago to have direct (fiber optic) connectivity with an AS in Seattle, and vice versa. Since the Zayo Group owns and operates the Zayo AS, zColo facilities and the fiber intercity connectivity, it is reasonable to assume that this fiber carries IP traffic of Zayo AS (AS6461) between the Chicago and Seattle colos. Zayo, additionally, provides detailed network maps (in KML format)[9] which contains extensive details on their facilities and intercity connections. Manual inspection of the maps revealed that zSea and zChi are indeed the endpoints of the aforementioned dark fiber between Seattle and Chicago. Interestingly, the fiber traverses another Zayo facility in Chicago, *i.e.*, the chain of colo facilities that the dark fiber connects is *zSea* → *zChi*2 → *zChi*. This fiber connectivity provides us with a unique opportunity for finding reliable anchors as follows. Using Zayo's traceroute LGs in Seattle and Chicago, we run traceroute probes from one colo to target IPs in Zayo at the other city in both directions. Our intuition is that these traceroute probes should traverse the routers in zSea and zChi and reveal the PoP-tags associated with them. Following this approach, we concluded that the PoP-tag associated with routers in zSea and zChi are `AS6461:ZAYO.SEA1` and `AS6461:ZAYO.ORD7`, respectively. Interfaces associated with these PoP-tags can, therefore, be used as anchors in the pinning process. We emphasize that an anchor does not need to be a physical entity

---

[8] An example of a graph including router nodes (representing IP aliases) and PoP-tag nodes is available at `onrg.cs.uoregon.edu/Impact/vis-js/net_pop_chi_aux_pinned_delay.html`

[9] `www.zayo.com/network/file-downloads`

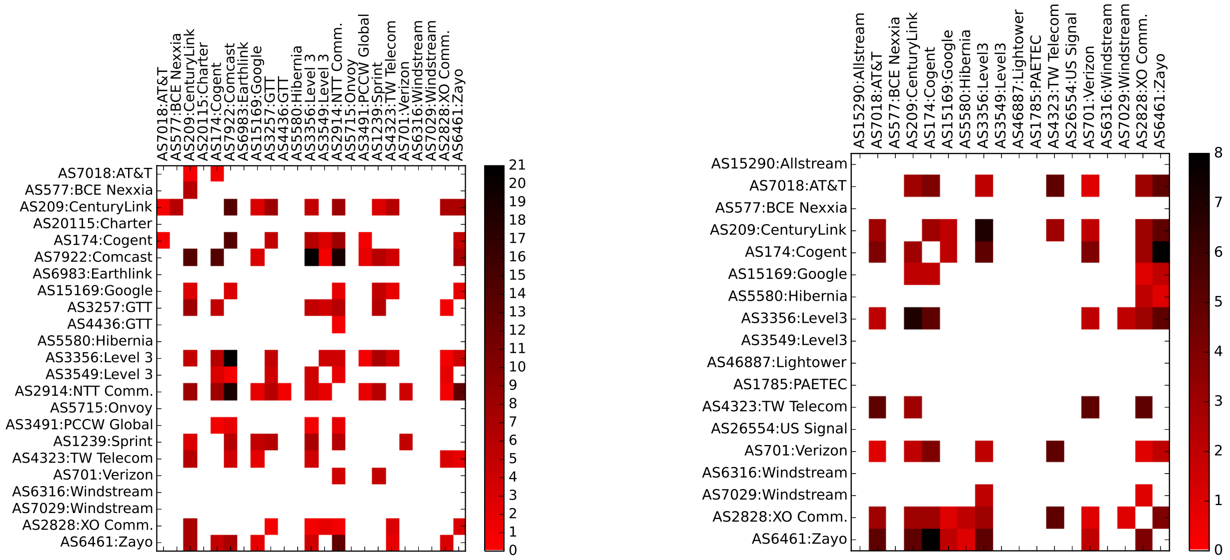

(a) zSea

(b) zChi

Figure 4: The number of physical connections identified in the target city between each pair of tenant ASes

(*e.g.*, an interface) at a target colo. Rather any attribute that can be reliably used to localize an interface to a particular colo, can be used as an anchor in the pinning process.

Using Regional ASes: The existence of regional ASes at a colo can also help in identifying anchors at that facility or validating already discovered anchors. Our assumption is that a regional tenant AS in a colo facility is likely to have a single PoP in that city where it establishes all its cross-connects. This assumption is supported by the economics of cross connects since being present at each colo facility has a cost that makes it prohibitive and unnecessary for a small regional ASes (with limited traffic) to be present at multiple colos. Instead, they are most likely to join a single colo whose tenants better match their needs. While we do not use this technique in this paper, it is often very useful in colo facilities with multiple regional ASes.

7.2 Inferring Colocation based on Association

Given a set of IP interfaces associated with the city-centric view of inter-AS links between tenant ASes along with some anchors for a target colo, our goal is to pin these interfaces (and thus map the links between them) to the colo facility. We start by initializing the label of all the interfaces to "close call" since they are all mapped to the target city. Then, we apply a few rules to identify a subset of these interfaces that are located at the colo facility (and change their label to "hit") and identify another non-overlapping subset of interfaces that are mapped to other facilities in the city (and change their label to "out"). Any remaining interface with a "close call" label may still be located in the target colo but we do not have sufficient evidence to make that inference.

Exclusion Rules: Before we pin any interface, we filter out interfaces that are highly unlikely to be at the target colo facility (and change their label to "out") by applying the following rules. The letter inside the parenthesis indicates the level of confidence–High (H), or Medium (M), or Low (L)–associated with the rule.

- *ExcOutIXP (H)*: Interfaces whose IP addresses belong to a known IXP-prefix, with the IXP located at a different facility, cannot be at the target colo.
- *ExcOutDNS (H)*: Interfaces with DNS names (PTR records) that indicate presence at a colo facility owned by a different company (*e.g.*, `eqnx`, or `equinix`, or `eqix` indicates that the interface is present at a colo facility owned by Equinix) cannot be in zColo.
- *ExcOutT1 (H)*: Interfaces of Tier-1 ASes (*i.e.*, pesudo tenants) that are unlikely to be physically in zColo since they typically house their PoPs in their own buildings as we described in Section 3.1.

Association/Inclusion Rules: We commence by pinning all interfaces that meet the anchor criteria, viz., containing PoP-tags identified as anchors in the Section 7.1, to the target colo facility. Then we apply the following heuristic rules in order to pin any remaining interfaces with close call label (the confidence associated with each rule is indicated within parenthesis).

- *Alias Association (H)*: Aliases of a router are collocated; if one of them is already pinned then the rest are pinned to that location.
- *PoP Association (M)*: Interfaces that have the same PoP-tag are collocated.
- *PoP Exclusion (M)*: Two different PoP-tags for the same

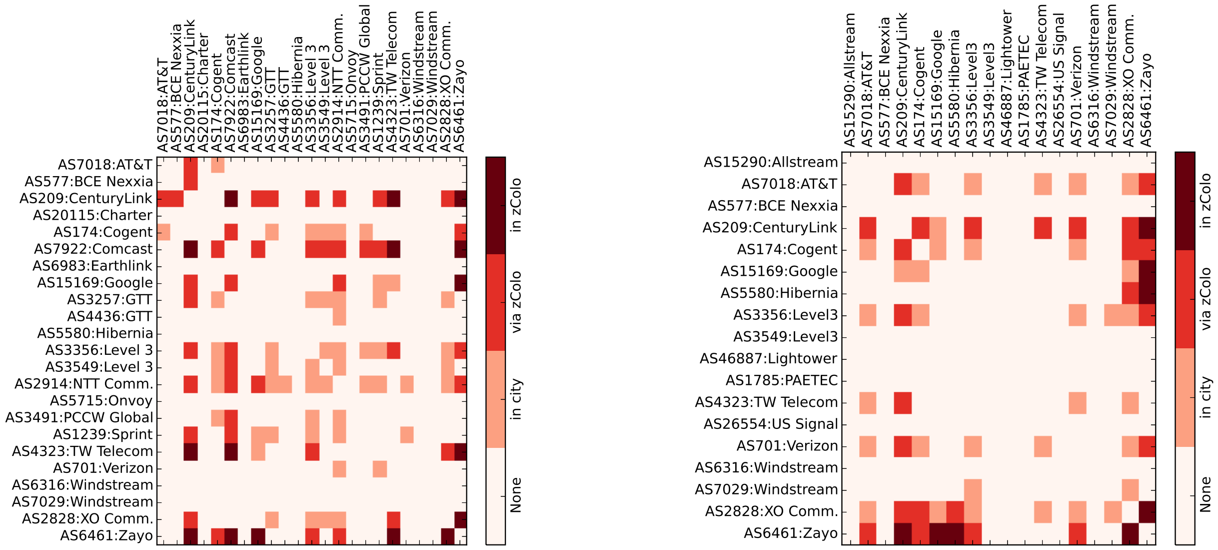

(a) zSea

(b) zChi

Figure 5: Pinning Status of all inter-AS links with respect to the zColo facility in each city.

AS within the same city (*e.g.*, `AS6461:ZAYO.ORD2` and `AS6461:ZAYO.ORD7`) *are not* in the same colo facility since each AS uses a distinct PoP-tag for each facility. Therefore, if one of the PoP-tags of an AS is pinned to a colo facility, the other PoP-tags of the same AS should not be localized to the same facility.

- *Delay Association (L)*: Interfaces at either end of inter-AS links are collocated *if* the delay between them is really short (*i.e.*, a couple of ms). We refer to this association as a *propagation rule* since it is the only rule that directly leverages the connectivity between interfaces for pinning.

These association rules are applied in two rounds as follows: In the first round, *(i)* we apply the first two rules to pin all interfaces with close-call label that are associated with anchors and change their label to hit. When these two rules lead to a conflicting label for an interface, the outcome of the alias association is selected as it is more reliable due to the potential DNS misnaming of IP interfaces [59]. *(ii)* Given the pop-tags of pinned interfaces, we can identify the PoP-tags that are associated with other colos and use them to apply the PoP Exclusion rule, mapping any close call interface to other facilities and set their label to out. *(iii)* We consider all the inter-AS links with short delays that one of their end interfaces is labeled as hit and apply the propagation rule to these links. This could lead to the change of label for a group of interface s due to pinning through propagation (PTP). In the second round, we only reapply the first three rules to pin any other interfaces that is associated with PTP interfaces. We conservatively apply the propagation rule only once since its iterative propagation could potentially lead to error. Note that once the label of an interface sets to hit or out, it does not change any further.

7.3 Mapping Results

We apply our pinning procedure to all interfaces associated with the city-centric view of inter-AS connections among tenants in zSea and zChi. Table 6 summarizes the main outcome of our pinning strategy. The top part of the table shows the number of pinned interface along with the corresponding mapped IP links and AS links at each colo. The bottom part of the table shows the break down of the number of the interfaces that are labeled as hit or out by individual rules for each colo. For alias and PoP-tag rules, we also show the number of pinned interfaces in each round (*i.e.*, before or after applying the delay association). In zSea, the 28 pinned IP links are associated with 8 AS links (one-third are P2P) connecting 5 ASes that consists of 1 content, 2 cable, and 3 T2. In zChi, the 10 pinned IP links are representing 4 AS links (half are P2P) between 5 ASes with this break down: 1 content, 1 regional, 1 cable, and 2 T2.

We also examine the sensitivity of our pinning procedure to the delay threshold for the propagation rule. Surprisingly, the value of this threshold does not have any effect on the outcome of our pinning procedure. Closer examinations revealed that the delay associated to all qualified inter-AS links for the propagation rule is less than 1ms. Therefore, any conservative delay threshold would work.

The assigned labels by our pinning procedure to each interface associated with inter-AS links in the target city leads to three possible pinning outcomes for links that are, from most to least desirable, as follows: *(i) in zColo* for links whose both sides are pinned to the colo facility, *(ii) via zColo* for links that only one side of them is pinned to zColo and the other side is either out or close call, *(iii) in city* for links

Table 6: Count of interfaces pinned by different rules

		zSea	*zChi*
	IP nodes	74	29
	IP links	28	10
	AS links	8	4
Rule	Alias Inc.	0+1	0+2
	PoP-tag Inc.	12+51	13+8
	Delay Inc.	11	7
	ExcOutT1	127	62
	PoP-tag Exc.	0+5	1+2

that both ends are labeled as out or close call. Each element in the symmetric matrix in Figure 5 provides a detailed view of mapped links by showing the most desirable pinning outcome among all the mapped IP links (if there is more than one) between a pair of tenant ASes. In summary, our pinning procedure mapped 8 out 33 (4 out of 33) AS links and 28 from 276 (10 out 109) IP links in zSea (zChi).

AS-Level Structure at zSea: Figure 6 visualizes the final outcome of our colo mapping effort for zSea as a graph (a similar figure for zChi is available in our technical report [40]). This graph captures the zSea-centric AS-level structure as well as interface-level connectivity. Each node indicates a tenant AS and each edge between two tenant ASes is annotated by the number of discovered interface-level connections between the pair of tenants. We use dark blue to show the ASes with at least one interface pinned to zSea. Light green shows Tier-1 ASes that are localized to another close-by facility and yellow is used for the remaining ASes. Cross connects that are mapped to zSea are shown as thick edges, and cross connects that result from connectivity options provided by zSea and are localized to a close-by colo are indicated by dashed lines. For the remaining edges, we have insufficient evidence to localize them to the inside or outside of zSea.

We also shared the final outcome of our mapping effort with the colo provider and the tenants in zSea and zChi and asked for feedback in the form of a questionnaire with the hope to be able to validate our findings. Unfortunately, we received less than a handful of responses, but those who responded confirmed the accuracy of our inferred cross connects for their networks in zSea or zChi. However, those confirmed cross connects made up only some 10% of all cross connects we placed in zSea and zChi, and our attempts to obtain any information from zColo were unsuccessful. An interactive version of the connectivity graphs that capture all or localized inter-AS links for both colo facilities are available online [10].

8. CONCLUSION

This paper proposes a methodology for localizing the cross connects that the tenants of a given colocation facility purchase from the colo provider to interconnect their networks.

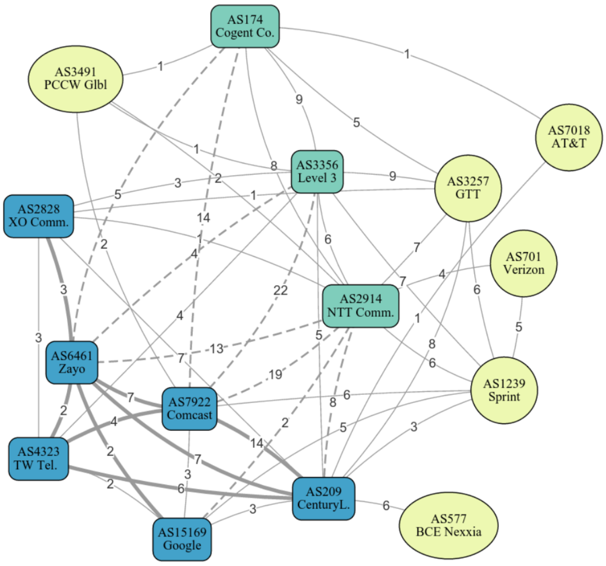

Figure 6: The visualization of mapped crossed connects between tenant ASes in zSea.

We demonstrate the feasibility of our approach by applying it to two medium-sized zColo facilities, one in Seattle and one in Chicago, and point out the many challenges that we encounter in this effort. In particular, we elaborate on the variety of interconnectivity options offered by different colocation providers and show how such "details" impact the task of determining the street-level geography of cross connects.

Our work is largely complementary to recent efforts aimed at accurately mapping the physical Internet. In particular, while [26] is concerned with constructing a map that shows the geography of the long-haul U.S. fiber-optic infrastructure (*i.e.*, constructing a network where nodes represent the different metro areas in the continental U.S. and links represent conduits that house the long-haul fiber-optic cables that run between them), [26] is not attempting to map the "insides" of the nodes, *i.e.*, it does not attempt to map the locations of all colo facilities and data centers in a given metro area or the fiber between them or the connections within them. In contrast, our study focuses exclusively on the "inside" part of these nodes and is not concerned with mapping the network's edges. The combination of the two approaches could provide an unprecedented view of the physical Internet, yielding a detailed account of the geography of both the long-haul fiber-optic cables that are used to connect the metro areas and the cross connects that are established inside the colos to interconnect different networks.

In future work, we hope to leverage edge-based measurement platforms such as Dasu [50] to expand our options for selecting suitable vantage points and targets, which may more often enable us to issue a traceroute that traverses a given router, PoP, or cross connect. At the same time, we plan to use the experience gained from mapping the two medium-sized zColo facilities when applying our approach

---

[10] `onrg.cs.uoregon.edu/Impact/xconnect/vis/`

to the 100+ large Equinix-owned and operated colos around the world where a large portion of all established cross connects are located.

9. REFERENCES

[1] Colocation cross connect prices vary greatly between US, Europe. https://goo.gl/mR3iFx.
[2] CoreSite Cross Connects. http://goo.gl/lAzjR8.
[3] CoreSite: Investor Presentation, February 2014. http://goo.gl/VgKoPQ.
[4] Equinix Q1 2015 Earnings Conference Call. http://goo.gl/IsJVXM.
[5] Equinix Reports Fourth Quarter and Year-End 2014 Results. http://goo.gl/yPgF4C.
[6] Long Haul Dark Fiber. http://goo.gl/TMgnBR.
[7] PlanetLab. http://goo.gl/bexM0.
[8] RIPE RIS. https://goo.gl/jFAsIA.
[9] Telx Cross Connects. http://goo.gl/IIPyk3.
[10] ZAYO Data Center Connectivity. http://goo.gl/T37cgO.
[11] Advanced Network Technology Center. University of Oregon Route Views Project. http://www.routeviews.org, 2013.
[12] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *SIGCOMM*, 2012.
[13] Akamai. EdgeScape Service Description. http://goo.gl/akSvGu, March 2002.
[14] A. Alur, R. Motamedi, and R. Rejaie. Inferring Attributes and Geographical Information from DNS Names. Technical report, Univ. Oregon, UO CIS, 2015.
[15] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding Traceroute Anomalies with Paris Traceroute. In *IMC*, pages 153–158, 2006.
[16] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *IMC*, pages 336–349. ACM, 2009.
[17] CAIDA. Archipelago (Ark) Measurement Infrastructure. http://goo.gl/nkxvgY.
[18] CAIDA. Skitter. http://goo.gl/qlDHVg.
[19] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is More to IXPs Than Meets the Eye. *SIGCOMM CCR*, 43(5), Nov. 2013.
[20] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. Quo Vadis Open-IX? *SIGCOMM CCR*, 45(1), Jan. 2015.
[21] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users. In *CoNEXT*, 2009.
[22] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: The AS-level connectivity Observatory. *SIGCOMM CCR*, 38(5):5–16, 2008.
[23] Data Center Map. http://goo.gl/v2F43, 2015.
[24] B. Donnet and T. Friedman. Internet topology discovery: a survey. *Communications Surveys & Tutorials, IEEE*, 9(4):56–69, 2007.
[25] R. Durairajan, J. Sommers, and P. Barford. Layer 1-Informed Internet Topology Measurement. In *IMC*, 2014.
[26] R. Durairajan, J. Sommers, W. Willinger, and P. Barford. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *SIGCOMM*, 2015.
[27] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. DomainImpute: Inferring unseen components in the Internet. In *INFOCOM*, 2011.
[28] D. Fitzgerald. Equinix Inc., the Internet's Biggest Landlord. http://goo.gl/eyGAq2.
[29] V. Giotsas, S. Zhou, M. Luckie, et al. Inferring multilateral peering. In *CoNEXT*, 2013.
[30] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier. Network topologies: inference, modeling, and generation. *Communications Surveys & Tutorials, IEEE*, 10(2), 2008.
[31] B. Huffaker, M. Fomenkov, and k. claffy. DRoP:DNS-based Router Positioning. *SIGCOMM CCR*, 44(3), Jul 2014.
[32] IP2Location. IP2Location, 2015. http://www.ip2location.com/, 2015.
[33] V. Katsenelson. How I Bought the Internet - and You Can Too. http://goo.gl/nB0n6M.
[34] K. Keys. iffinder. http://goo.gl/4BFqWn.
[35] M. LLC. GeoIP, 2010. http://www.maxmind.com, 2010.
[36] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, et al. Challenges in Inferring Internet Interdomain Congestion. In *IMC*, 2014.
[37] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *OSDI*, pages 367–380, 2006.
[38] P. Marchetta, V. Persico, E. Katz-Bassett, and A. Pescapé. Dont trust traceroute (completely). In *ACM CoNEXT Student workshop*, 2013.
[39] M. Mooney. "Chicken" — A Game Played as a Child and by some ISPs with the Internet. http://goo.gl/tzx8D9.
[40] R. Motamedi, B. Chandrasekaran, B. Maggs, R. Rejaie, and W. Willinger. On the Geography of X-Connects. Technical report, Univ. Oregon, UO CIS, 2015.
[41] R. Motamedi, R. Rejaie, and W. Willinger. A Survey of Techniques for Internet Topology Discovery. *Communications Surveys & Tutorials, IEEE*, 2015.
[42] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (in)completeness of the observed Internet AS-level structure. *IEEE/ACM Transactions on Networking (ToN)*, 18(1), 2010.
[43] Packet Clearing House. Routing archive. https://goo.gl/63xfy8.
[44] J.-J. Pansiot and D. Grad. On routes and multicast trees in the Internet. *SIGCOMM CCR*, 28(1), 1998.
[45] PeeringDB. Exchange Points List. https://goo.gl/urdmmR.
[46] O. D. Project. Colocation. http://goo.gl/byPSwO.
[47] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *IMC*, 2014.
[48] N. RIPE. Ripe atlas. ttps://atlas.ripe.net, 2010.
[49] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 Lessons from 10 Years of Measuring and Modeling the Internets Autonomous Systems. *Selected Areas in Communications, IEEE Journal on*, 29(9), 2011.
[50] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing experiments to the Internet's edge. In *Proc. of USENIX NSDI*, 2013.
[51] Y. Shavitt and N. Zilberman. A geolocation databases study. *Selected Areas in Communications, IEEE Journal on*, 29(10), 2011.
[52] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *SIGCOMM*, 2002.
[53] M. Taylor. Verizon's Accidental Mea Culpa. http://goo.gl/NvyMLl.
[54] Team Cymru. IP to ASN Mapping. https://goo.gl/Sb2oQh.
[55] UCLA Internet Research Lab. Internet AS-level Topology Archive. http://goo.gl/2SOJF, 2015.
[56] K. Xu, Z. Duan, Z.-L. Zhang, and J. Chandrashekar. On Properties of Internet Exchange Points and Their Impact on AS Topology and Relationship. In *Networking*. 2004.
[57] D. Young. Thoughts on Internet congestion and the FCC's broadband report. http://goo.gl/Gc3Re6.
[58] D. Young. Why is Netflix Buffering? Dispelling the Congestion Myth. http://goo.gl/1VMTi2.
[59] M. Zhang, Y. Ruan, V. S. Pai, and J. Rexford. How DNS Misnaming Distorts Internet Topology Mapping. In *USENIX ATC, General Track*, 2006.