

# Framing the Value of Internet Exchange Participation

Jesse H. Sowell

**Abstract**—Internet eXchanges (IXes) are developed to localize traffic, reduce connectivity costs, and reduce latency. Historically, transit savings were sufficiently substantive to justify investment in IX interconnection strategies, overshadowing other benefits enjoyed by IX participants. Recent work argues these basic IX effects facilitate regional economic development. In developed regions, as transit prices converge with comparable per volume price on Internet exchanges, a new question emerges: given this convergence, what additional value and benefits does the IX provide existing and would-be participants? Governance and organization studies of IXes provide some of the nuance necessary to answer this question. Based on interviews with IX managers and participants, there is value in a more fluid market of interconnection options that facilitate investment deferral and greater control over the repurposing of interconnection resources. Interconnection is presented in terms of platforms and how they implement interconnection options, highlighting how the mechanics of these implementation frame network actors' decisions regarding how to maximize access to and investments in interconnection markets.

Semi-structured interviews of a wide range of network operators and IX operators over the last year and a half imply decision heuristics for optimizing value and network resilience: selectively increasing redundancy, increasing unique interconnection partners, and reducing switching costs, are among the most common criteria. This work distills decision heuristics identified thus far into a *partially parameterized model* of interconnection decision making across platforms that serve as markets for interconnection options. Unpacking decision heuristics contributes to more precisely explaining the mix of interconnection options available to actors that both derive value from network- and application-level services *but who also* influence infrastructure investment and strategy decisions through the options they select. This work demonstrates how IX-platforms facilitate greater access to interconnection options and defer bargaining and measurement costs necessary for more specific asset investments. IX-enabled interconnection contributes to the “flattening” of the Internet by providing an investment path to more sophisticated bundles of (flatter) bilateral relations rather than participating solely in the transit hierarchy. The model developed here provides a clear formalism for comparing these.

Two classes of parameters are missing: connectivity cost information and parameters representing actor-specific valuations of system properties such as redundancy or latency. A key contribution is to understand the options available by carefully *specifying* the parameter space based on empirical observation: what are the variables and how are they conceptually related? This work unpacks the mechanisms implementing two common types of interconnection option and concludes with working hypotheses that will be further refined and tested in ongoing work. The *next* stage of this work focuses on *identifying* appropriate operationalizations, data sources for cost, and strategies for comparing (*validating*) theoretical valuations based on empirical operationalizations of variables such as redundancy.

A social science approach to indicator development is applied to frame heuristics as background concepts that are refined into specifications of incentives, relationships, and potential indicators. Variables specified here represent possible systematizations of background concepts elicited from interviews. Neither the specification nor the relationships are presented as being writ in stone—rather, they are a part of an ongoing project to elicit empirical trends in the interconnection industry. The variety of network actors that engage with IXes, their value propositions, and the means by which they connect to IX platforms provide insight into which interconnection options are appealing to which types of actors as well as common trade-offs. Qualitatively, this gives insight into which types of actors are willing to invest in which mixes of infrastructure types and topologies. By highlighting trade-off conditions in the demand for interconnection, this specification provides direction for further identification of interconnection costs and the constraints by those making investment decisions.

The framework offers a number of contributions. Different interconnection modes are conceptualized independently, but the framework subsequently highlights mutual dependence and benefits. Working hypotheses argue that different platforms structure and mediate interconnection option implementations (transport, colocation, IX), how they are interleaved with one another in practice, and offer a first pass at identifying and realizing critical paths through investment decisions as they grow and develop specific relationships (and asset investments) beyond simple transit. Building on the notion of a real option, flexibility in general purpose resources (IX-mediated interconnection) facilitates staged, dynamic investment decisions and learning effects.



## 1 INTRODUCTION

In the late 1990's Internet eXchanges (IXes) were developed to improve performance by exchanging geo-

*This work has been funded by the Office of Naval Research under award number N00014-09-1-0597. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the Office of Naval Research. Ongoing work is funded by a Google Research Award.*

graphically local traffic in a local interconnection platform. In the nearly 25 years since, IXes have become key platforms in the Internet infrastructure ecosystems of the US, the EU, and other developed regions. Development efforts in regions of Africa, South American, and Asia aspire to replicate IX successes and the development benefits for the Internet industry writ large. Immediately observable and measurable outcomes in developed (and some developing) regions with

thriving IX ecosystems are reduced transit costs and lower latency between local IX participants. Although still compelling in many contexts, transit savings and latency are not the only benefits.

Interviews with IX participants and operators include to more nuanced benefits. On the IX, operators describe heuristics for greater resilience and better informed deployment and provisioning decisions. Within the infrastructure value network, IX and network operator interviews also claim IXes have contributed to *a*) erosion of transit prices; *b*) transport market development, in particular a transition of transport-IX relations from substitute to complement; *c*) encouraging the development and standardization of colocation facilities, and *d*) demand feedback from actors leveraging interconnection bundling strategies in markets such as content delivery and application performance.

This work builds on heuristics to comparatively frame interconnection bundling strategies across interconnection platforms, in particular the varieties of co-evolution of IX and colocation platforms. Conditioned on the value-proposition of a given actor, formalizing these heuristics provides a framework for reasoning about the boundary conditions of different interconnection strategies, often highlighting a mix of different modes of interconnection provisioning. This framework is presented as the first step in a larger model for context-based evaluation of the demand-side value, and viability of, IX-mediated interconnection platforms.

Mature Internet infrastructure markets such as the EU enjoy the *combined* benefits of diverse interconnection provisioning options. Not surprisingly, this is due in large part to competitive transport and colocation markets. In contrast, many developing regions are dominated by incumbent carriers, having limited transport and colocation markets: in effect they are limited to carrier-facilitated transit. IXes have been offered as means to catalyze Internet infrastructure development in developing regions. The framework presented here captures some of these constraints on interconnection bundles, and by proxy, available interconnection strategies. The larger work aspires to leverage comparisons across regions to develop contingent theories of how interconnection provisioning investment can be incentivized.

The most intuitive and historically most commonly cited benefit of IX-mediated connectivity is the opportunity to reduce transit costs. This is appealing to networks in developed and developing regions alike. Figure 1 depicts a stylized view of per volume costs of IX ports relative to transit savings in Europe. The right panel of Figure 1 highlights transit price erosion, convergence with IX costs, and the supposed decline in overall IX value. When IXes were emerging in the late 1990's, many of the founding networks were ISPs whose primary cost was in fact transit—coupled with a period in the industry characterized by high risk and

low margin, transit savings was a substantive benefit in and of itself.

Today, Figure 1 tells a seemingly different story. Transit erosion is in part a function of Europe's competitive transport and colocation markets and in part a function of an IX ecosystem that reduces the transaction costs of interconnecting at colocation facilities. A trite, albeit somewhat accurate, explanation is that IX efforts at reducing transit costs may have worked too well: now they face price convergence and according to some, market saturation. The skeptical network peering coordinator may well say: "If IX connectivity is only marginally less expensive than transit *and* the cost is continuing to converge with transit, what are the benefits when I still have to establish some form of transit anyway?"

Interviews and community investment in IXes tell a different story than the skeptical framing above. Participating in the IX ecosystem does have value beyond transit savings, but, as hypothesized here, is strongly predicated on the specific-value proposition of the IX participant. For instance, the value-proposition of CDNs, cloud application providers, and gaming networks rely on low latency and enhanced topology insight garnered from IX participation. Eyeball networks want one-stop shopping for local and remote content, ideally over peering links. Network architects argue that more diverse views of routes, prefix advertisements, and potential interconnection partners facilitate resilience and better planning. For instance, when leveraging multiple interconnection platforms, networks can develop interconnection strategies around redundancy, finding unique interconnection partners, or both. The options framework developed here lays the foundation for addressing the skeptical peering manager's question, highlighting the information available to make investment decisions. The partially parameterized model highlights what additional data is needed to differentiate context and value-proposition specific critical paths through the option trade-off space faced by the skeptical peering manager.

A key premise is that the value of, and subsequently the demand for, collateral benefits to IX-mediated interconnection is *actor- and context-specific*. Two key objectives are articulated across the board: the ability to reason about the availability of unique interconnection options and the opportunity to invest in redundancy. Each is characteristically different in terms of investment and deferral options. Availability of these options is partially a function of interconnection provisioning. Key questions are:

- 1) How are interconnection relationships provisioned?
- 2) What role do bargaining and measurement costs associated with this decision play in resource investment?
- 3) What kinds of information are available for strategic interconnection bundling?




Figure 1. In Europe, a number of IXes are faced with transit price erosion and convergence with the costs of traffic in the IX.

An options framework illustrates the dynamic decision-making processes involved in interconnection investment relative to an existing interconnection bundle) based on differences in provisioning and related costs. A hypothesis that will be developed over the course of this paper is that for certain actors, the diversity and flexibility afforded by IX-mediated interconnection options is a cost effective option that allows actors to defer the decision to invest in colocation or transport-mediated interconnection until sufficient information regarding new relationships becomes available.

## 2 CONTRACTING MODES AND INTERCONNECTION PROVISIONING

Interconnection contracts and their economic implications have been discussed in the literature, most notably by Faratin et al. (2007) and Clark, Lehr, and Bauer (2011). Varieties of transit and peering are presented as contracting modes dictating traffic ratios and costs once interconnection is established. Faratin et al. (2007) describes strategic implications by comparing static bundles of interconnection relationships. Tacit in this analysis is the availability of different contracting modes. The framework presented here develops the role of IXes as a low-cost means explore the diversity of contract bundles that may be established. The options model highlights how framing this as a dynamic decision process can frame the feasibility and information necessary to develop strategic interconnection bundles. Different modes of interconnection have different barriers, including capital provisioning costs, bargaining costs, and measurement costs. These modes are also mediated by different actors with different cost structures of their own.

Understanding the character of access to interconnection options and the role of IXes requires reasoning about the viability of exercising an option on different platforms. Different models of interconnection *provisioning* differentially structure availability and asset specificity. *Establishing* interconnection comprises to distinct activities that will be further unpacked: provisioning of the physical link(s) between network actors  $i$  and  $j$  and establishing a contractual relation

between  $i$  and  $j$ . In this work, order matters and is a function of investment decision deferral. Network actor  $i$  and  $j$ 's<sup>1</sup> decision may be based on limited or extensive information about how much traffic they exchange. This work assumes that different interconnection configurations, capacity, and latency will affect traffic levels.<sup>2</sup> Under this notion of interconnection, an *interconnection relation* means that network actors  $i$  and  $j$  have both *a*) provisioned, or contracted the provisioning of, the physical link(s) necessary to establish interconnection *and b*) have establishes traffic ratios, cost settlements, and other aspects of the contractual component of the agreement.<sup>3</sup> As illustrated in Figure 2, elements of provisioning and contractual negotiation are interleaved in various ways depending on the platform mediating interconnection. A key goal in selecting platform bundles is to minimize bargaining and measurement costs while balancing general and specific asset investments. Here, that trade-off space is structured by interconnection platforms.

Interconnection platforms serve as a loci of interconnection potential, framed here as interconnection marketplaces in the larger interconnection market. The notion of an interconnection option has been used colloquially thus far.. An *instance* of an *interconnection option* on a given platform  $p$  is the *option*, *but not the obligation*, for  $i$  to interconnect with any other network actor  $j$  that also participates in platform  $p$ . When  $i$

1. For consistency throughout this paper considers interconnection from the perspective of network actor  $i$ . The interconnection partner on the other side of the relation is typically labeled  $j$  or some  $j_i$  in a larger set of potential interconnection partners  $1 \leq i \leq n$ .

2. It is naive to assume some mode of static demand. Rather, interviews have highlighted that when quality of experience increases, demand increases. For instance, when a CDN cache is introduced the demand for those services increases. The nominal set of services has not changed, but the quality of experience has, tapping what may be considered latent demand. While it is well known amongst experienced operators in the community that latent demand exists and manifests, the magnitude is often unclear and surprising. It is this uncertainty that makes specific asset investment difficult and warrants engaging in a low-cost investment sufficient to “sample” the growth pattern before making a more specific asset.

3. The term contractual relation does not necessarily imply a formal written contract. It may be a handshake based on well-known convention. That said, one aspect of future surveys will be to determine how these conventions have changed. This will in part replicate Woodcock's study.

provisions connectivity to a platform, it is exercising a *platform option* in order to establish access to the options available at platform  $p$ .

Interconnection options are further refined by elaborating how different interconnection platforms provision links amongst their participants. This framing highlights key empirical questions:

- 1) How do the cost structures differ?
- 2) How do these differ in terms of bargaining and measurement costs?
- 3) What are the implications of these differences for the deferral of investment decisions?
- 4) What are the implications for developing strategic interconnection bundles?

The two option provisioning strategies described here differ in terms of timing, constraints, cost division, and platform infrastructure investment.

These distinctions set the stage for discussing the effects of options' availability on strategic bundling of interconnection relations *and* options based on the availability options, perceived viability, and decision making processes facilitated by the different platforms. The focus in Section 3 is on the mechanics of availability; heuristics used to evaluate viability are discussed in terms of comparative bundles in Section 4. To bridge the gap between mechanics and the interconnection relation as a whole, the next section reviews the contractual component of interconnection.

## 2.1 Contractual Relations: Transit and Peering

Faratin et al. (2007) discuss interconnection bargaining games in terms of fairly static interconnection bundles. These games highlight the shift from simple contracting modes like transit and peering (discussed shortly) to the rise of partial transit, paid peering, and, most importantly, strategic permutations of these used to leverage market power. Faratin et al. (2007) also highlights that the actors have changed since canonical notions of transit and peering first evolved—ISPs are but one of many different actors in the Internet infrastructure ecosystem. Moreover, asymmetries between actors give rise to strategic interconnection practices. Sophisticated content delivery architectures and large eyeball networks represent a contemporary value network whose business negotiations play out in part through strategic manipulation of bundles of interconnection relations in the larger interconnection market, across a variety of platforms.

Clark et al. (2011) contemplates the competitiveness of interconnection markets and the role of transparency in contractual relations on how these markets work. Both works address strategies that shape contracting modes of interconnection relations by comparing static bundles and their implications for market power. This work complements those discussions by explicitly framing bundles as sets of relations

whose development and adaptation to strategic objectives incur bargaining and measurement costs.<sup>4</sup> More specifically, different interconnection option implementations structure bargaining and measurement costs differently: one hypothesis is that IXes provide a greater proportion and diversity of actors in the Internet infrastructure ecosystem with the opportunity to develop more sophisticated interconnection bundles. Later, this is framed as “in-sourcing” strategically beneficial interconnection relations previously relegated to transit (see Section 4.2.4). To reason about these bundles, a background in the contractual relations therein are necessary. The following sections provide a primer on transport as a connectivity option, followed by discussions the varieties of transit and peering that are at play in the interconnection market.

### 2.1.1 Transport

*Transport* provides a point-to-point (or circuit) from the point of presence (POP) of one facility ( $i$ ) to the POP of some other facility  $f$ , which may be a POP of  $j$  or an interconnection platform  $p$ . Network actor  $i$  has a variety of options to establish transport. One case is for  $i$  to build its own physical connectivity to  $j$ . Building infrastructure de novo is expensive and unlikely unless  $i$  is itself a transport or transit provider. Some of  $i$ 's typical options are: *a*) lease existing physical infrastructure, for instance leasing dark fiber  $i$  can light itself; *b*) purchase a set of wavelengths in existing fiber; *c*) contract transport services from a provider  $c$  that has more extensive infrastructure; *d*) contract transport platform services, such as IX-Reach or Atrato that have specialized in facilitating connectivity to interconnection platforms. More exotic transport options include transport over power lines and long distance wireless communication. Both have seen deployment in underdeveloped and/or geographically challenging regions. For instance, wireless has seen use in mountainous regions, hopping from ridge to ridge or in locations with little terrestrial infrastructure.

4. Two bargaining games overlap here. The bargaining discussions in Faratin et al. (2007) and in Clark et al. (2011) are both from the perspective of existing bundles. In particular, the discussion of the flow of content from Level3 to Comcast revolved around traffic rations in an existing interconnection relation and who should bear the cost of capacity upgrades. One way to frame this is that the interconnection relation is a shared resource whose capacity is governed by the contractual relation. As will be discussed later, the bargaining costs in developing interconnection relations are presumed to require information regarding traffic ratios and will be differentiated in terms of whether this is a new relation for  $i$  and  $j$  or whether it updates an existing relation. For simplicity, these bargaining costs are scoped to the discussion of *mutually* planned interconnection capacity provisioning and upgrades. The Level3-Comcast dispute is an exception to mutual planning. The empirical question of when mutual planning breaks down or is infeasible is brought up in the concluding discussion. For the majority of the discussion here, value will be predicated on the *opportunity* for mutual planning and its exercise. That said, the conclusions recognize this assumption is not universal and briefly describes how it will be addressed in ongoing work.

In and of itself, transport does not guarantee an interconnection relation. Transport is one step in gaining access to a loci of interconnection. When establishing transport to the (dedicated) facilities of  $j$ , the only “option” available is to interconnect with  $j$ . In contrast, transport to an interconnection platform is the first step to gaining access to options with participants on that platform. In particular, one of the factors that makes these loci of interconnect attractive is the presences transit providers, discussed next.

### 2.1.2 Varieties of Transit

*Transit* means a network actor  $t$  ensures that those to whom it provides transit service have, by virtue of  $t$ 's (strategic) bundle of interconnection relations, connectivity to the rest of the Internet. A *simple transit relation* means network actor  $i$  has established an interconnection relation with  $t$  comprising *a*) physical interconnection at some facility  $f$ ,<sup>5</sup> and *b*) the contractual guarantee that the rest of the Internet is reachable by virtue of subsequent connectivity. In effect,  $t$  guarantees  $i$  access to the rest of the world. When  $i$ 's connectivity comprises of only transit from a single provider  $t$ , network  $i$  relies *exclusively* on  $t$  for all external connectivity. In terms of connectivity assets and decisions,  $i$  has exported all off-net connectivity to  $t$ , in particular geographically local routing decisions.<sup>6</sup> This is also referred to as *single homing*.

To get from  $i$  to  $t$ , underlying transport connectivity  $c$  may be provided by  $t$ , denoted  $c_t$ , or by third party, denoted  $c_n$ . Under the simple transit relation, if there is a failure in  $c$  that disconnects  $i$  from  $t$ , the result is a complete loss of connectivity for  $i$ . If there is a failure in  $t$  itself, the outage may be complete or partial. Historically, to mitigate connectivity outages, a network establishes multiple transit relations, typically a primary and a backup transit relation.<sup>7</sup> Generalizing, from the perspective of  $i$ , an  *$n$ -redundant transit bundle* is one in which  $i$  establishes  $n$  transit relations to ensure that if one fails, it has  $n - 1$  operational transit relations, each sufficient to reach the rest of

5. This is a simplification. Many network connect at a number of facilities. This is a simplification for clarity here, the constraint will be relaxed in later discussions.

6. This statement is not meant to imply  $i$  has control over all routing decisions affecting packets it sends into the Internet. Through strategic selection of interconnection partners it can influence routes and avoid others. The value of this level of control to  $i$  is a key empirical question; a coarse-grain proxy for that value is represented by  $K$  in Section 4.

7. Having multiple upstreams is known as multi-homing; one may multi-home for redundancy and load-balancing.

the Internet.<sup>8</sup>

A *partial transit relation* between  $i$  and provider  $t$  only provides  $i$  with connectivity to a subset of networks in the Internet. Network  $i$  may engage in multiple partial transit relations to create an interconnection bundle with the same effect as full transit (the entire Internet is reachable). Network  $i$  may also use partial transit from provider  $u$  to add redundancy to a distinguished subset of networks covered by existing transit  $t$ . Such a set of interconnection relations will be referred to as a *mixed transit bundle*. Following the market metaphor, interconnection platforms comprising transit providers, are a natural market for transit relations, arguably fostering competition due to low switching costs relative to building to the dedicated facilities of a transit provider. Interconnection platforms are also a market for other contracting modes, such as varieties of peering.

### 2.1.3 Varieties of Peering

*Peering* is a mode of interconnection in which networks  $i$  and  $j$  agree to exchange traffic originating in their networks and their downstreams<sup>9</sup> at a given point of interconnection. As above, networks may peer directly if there is shared transport  $c$  between dedicated facilities of the two networks; in effect  $c$  provisions the link necessary for the interconnection relation. This kind of interconnection is a specific asset investment. For this analysis, it is assumed  $i$  and  $j$  have a well-understood bilateral relationship that warrants investment in directly connecting facilities. Commonly, and the focus of this work,  $i$  and  $j$  establish connectivity to common interconnection platform and interconnect there.

Peering is nominally defined in terms of the origins and destination of traffic and the ratios of traffic. The exchange of payment, if any, is typically governed by the ratios of traffic between  $i$  and  $j$  over the peering link.<sup>10</sup> In *settlement-free peering* both parties agree that the traffic ratio between  $i$  and  $j$  is sufficiently close to equal ( $i \approx j$ ) that the transaction cost of measuring the difference is higher than the value of the difference to either  $i$  or  $j$ . For example,  $i$  may technically send

8. Ideally, each of the  $n$  transit relations has independent transport connectivity  $c_j$ , where  $j = 1 \dots n$ , such that a failure in transport only disconnects one transit relation. Alternately, consider an  $n$ -redundant transit relation where a single connectivity provider  $c$  is the transport supporting all  $t_{1 \dots n}$  transit relations. Failure of  $c$  could result in loss of all connectivity, despite  $n$  redundant transit providers. Multiple transport relations  $c_{1 \dots n}$  reduce the risk of failure in overall connectivity; henceforth  $n$ -redundant transit relations will be assumed to have this property. This scenario protects against failures in connectivity to the transit network, but not failures of connectivity *in* the transit network.

9. Downstreams of  $i$  are networks that pay  $i$  for some form of transit. If a downstream network  $d$  pays for full transit,  $j$  will have access to  $d$  through  $i$ . If  $d$  pays for partial transit and  $i$  is not included in the subset covered by the partial transit agreement,  $j$  probably will not include  $d$  as a network reachable in its peering agreement with  $i$ .

10. In the case of larger networks, a peering relation may comprise many links but the limitation to establishing paths between  $i$  and  $j$ 's prefix cones remains.

more ( $i \gtrsim j$ ) but not enough to warrant monitoring and accounting. Historically this was also premised on the idea that  $i$  and  $j$  are approximately the same in terms of size and per unit of traffic operational costs. As per Faratin et al. (2007) and Clark et al. (2011), the heterogeneity of “ISPs,” here more generally referred to as network actors, means both network infrastructure investment and traffic levels vary substantively across network actors participating in common value networks that require interconnection.

*Settlement-based*, or *paid*, peering has the same technical characteristics in terms of the links necessary to interconnect and the origins and destinations of traffic. Paid peering differs in that one of  $i$  or  $j$  pays the other to engage in the peering relation. Consider an oversimplified case of traffic asymmetries. When the ratio of  $i : j$  is much higher than  $1 : 1$  ( $i \gg j$ ),  $i$  is sending much more traffic to  $j$  than  $j$  is sending to  $i$ . There may be a variety of reasons for this. A very simple interpretation is that, all other things being equal,  $i$  is *creating* greater amounts of traffic, and subsequently greater costs, for  $j$ . Under this framing,  $i$  should pay  $j$ .

Interconnection economics is not that simple. Consider the standard eyeballs and content argument. For instance, end-users in  $j$ 's prefix cone may place a high value on traffic from  $i$ . Under one point of view,  $j$  should extract some of that value from those end-users to compensate for the greater costs created by end-user demand. Under a different point of view, assuming  $i$ 's traffic is valuable and  $i$  is deriving value from either delivering to end-users or directly from end-user consumption,  $i$  should compensate  $j$  for the additional operational costs from that derived value. An interpretation of the former is that  $j$  is paying for content. An interpretation of the latter is that  $i$  is paying for access to end-users.

Yet a third interpretation is that neither  $i$  nor  $j$  want to have to risk user attrition by raising prices that cut into consumer surplus. Thus  $i$  and  $j$  vie for the interpretation (and bargaining position) that places the burden of extracting greater value from end-users on the other. This is a common and well-known problem of determining whether value should flow in the direction of content or counter to the direction of content. Refining this interpretation, one hypothesis is that there is not sufficient credible information about actual resource utilization and the attendant costs in particular interconnection value networks (such as content delivery) to negotiate demonstrably mutually beneficial settlement-based contracts.

This work develops the fourth interpretation by unpacking the means of establishing interconnection relations and the information available to actors when considering the viability of an interconnection bundle comprising both extant relations, unexercised options, and newly available options. As highlighted by Faratin et al. (2007) and Clark et al. (2011), different bun-

dles of relations reflect different strategies, different interpretations of the flow of value, and how these flows should be divvied up amongst those provisioning underlying resources. The availability of options is argued to increase the availability of these attendant strategies by introducing additional information into the decision processes. In some cases the perceptions may be entrenched, in others there may be opportunities to establish information sharing norms that facilitate coordinated deployment. Platforms' provisioning strategies, and consequently the availability of options, is discussed in the next section.

### 3 INTERCONNECTION PROVISIONING

Establishing interconnection between network actors  $i$  and  $j$  requires provisioning some means to connect their physical networks. Three means of establishing the links that comprising physical connectivity necessary for interconnection are presented: *a*) transport-mediate (or “dedicated”) bilateral interconnection; *b*) colocation-mediated interconnection (or simply colocation); *c*) IX-mediated interconnection. These differ in how interconnection *a*) is provisioned, *b*) by whom, *c*) the loci of control (mediation<sup>11</sup>) in that provisioning process, *d*) how easily the necessary resources are repurposed, *e*) and bargaining costs. Transport-mediated interconnection between  $i$  and  $j$ , discussed in the next section, is considered the baseline comparator. It is also most specific asset investment in interconnection considered here. The other two interconnection provisioning mechanisms are mediated by platform providers.

To illustrate and compare these processes, Figure 2 depicts the sequence of actions necessary to establish an interconnection relation under each provisioning regime. What is important in this diagram, elaborated in the next sections, is how and where in the sequence each platform mediates the process and the implications for the utilization and repurposing of resources in play. It is very important to note that this is a simple, generic illustration of what is considered a typical interconnection process; variants are an empirical question that will be explored further in ongoing work. Actual colocation facilities and IXes offer variants of these generic themes—this depiction is to illustrate the general process, not capture every variant in play.

#### 3.1 Transport-Mediated Provisioning

Under *transport-mediated provisioning*,  $i$  and  $j$  establish point-to-point transport between two dedicated facilities. Provision of transport presumes  $i$  and  $j$  have established a contractual relation,  $c_{neg}$  in Figure 2), part of which indicates how transport costs

11. Mediation here means that  $i$  and  $j$  must coordinate with the provider to establish interconnection and to contract capacity. Providers *do not* interfere with the flow of traffic beyond enforcing capacity constraints.

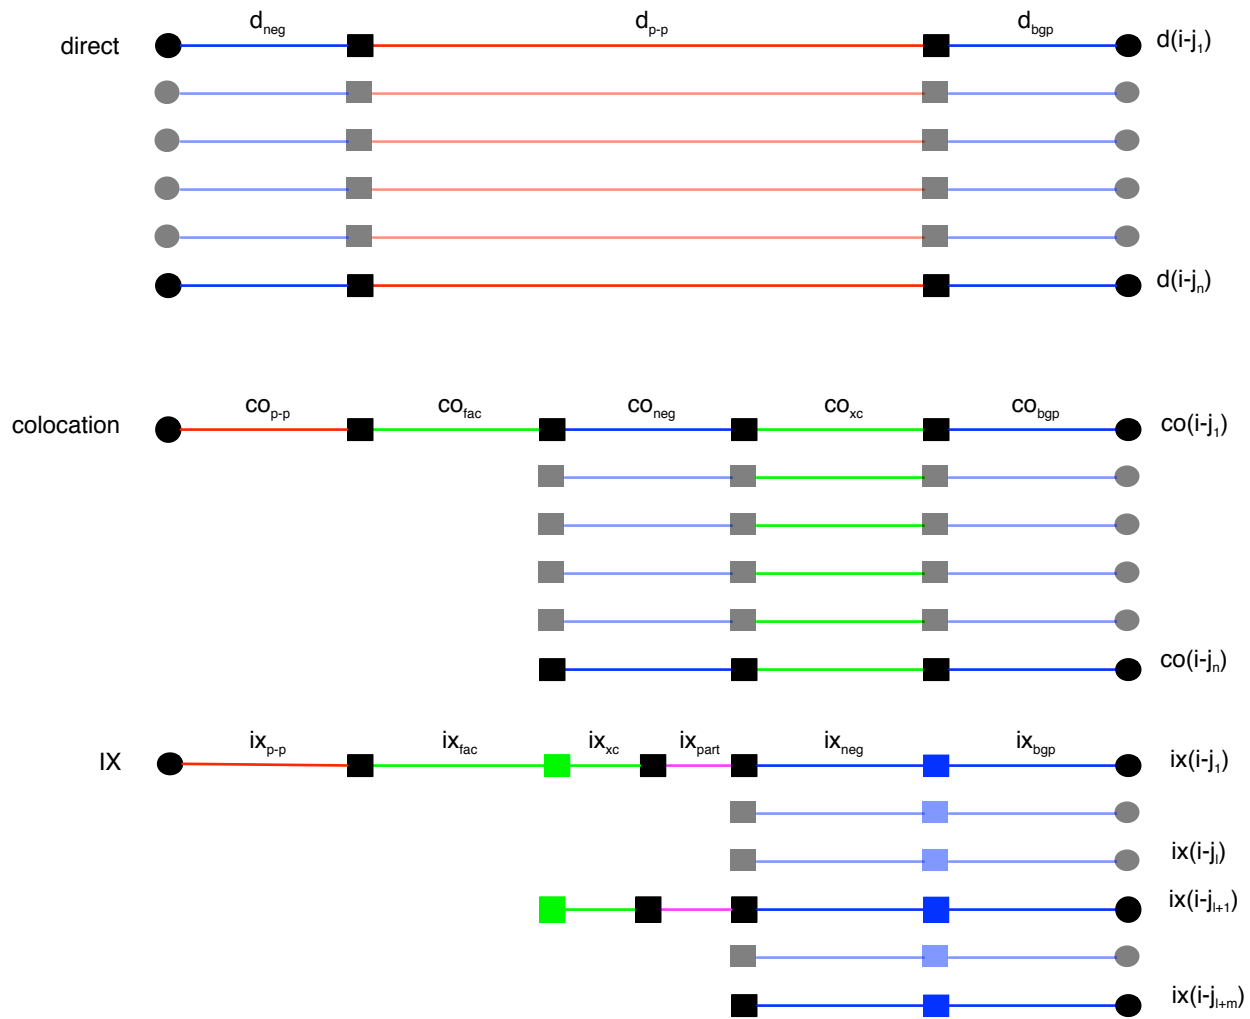


Figure 2. Three types of interconnection provisioning are presented: direct connections are denoted  $d$ , colocation  $co$ , and IX  $ix$ . Each horizontal line depicts the sequence of actions necessary to establish interconnection via that particular platform. Each step is denoted by the platform, subscripted by the type of action: for instance, establishing point-to-point connectivity to a colocation facility is denoted  $co_{p-p}$ . The actions are:  $neg$ , negotiation of interconnection contract terms;  $p-p$ , point-to-point transport;  $bgp$ , establishment of the BGP session between  $i$  and  $j$ ;  $fac$ , establishing residence at a hosting or colocation facility;  $xc$ , establishing a cross connect at a facility; and  $part$ , establishing participation on an IX. Colors denote who interacts in a given decision: red indicates interaction with a transport provider; green is interaction with a colocation facility; purple is an interaction with an IX; blue signifies interactions solely between  $i$  and  $j$ . Circles represent start and end of interconnection relation development; square delimit the different actions.

are allocated. Part of this presumption is that  $i$  and  $j$  understand the levels of traffic they exchange and have entered into a mutually beneficial contract. There are certainly corner conditions, but the important point within an options framework is that, for this framing of interconnection provisioning problems, a direct connection is presumed to be based on information justifying a substantive, relatively static investment. Further, given that it is a connection to a single network  $j$  and that  $j$  is not a loci of interconnection in and of itself, investment in this connection does not create opportunities for later dynamic decision making

that characterizes an option. In transport-mediated interconnection, there are no deferred interconnection options as with platforms; rather, it is a static investment decision. Under direct point-to-point provisioning between the facilities of  $i$  and  $j$ , there is only one choice available after the investment at each end of the transport provisioning investment: interconnect with the network at the other end.<sup>12</sup>

In Figure 2, each direct provisioning investment is

12. This assumes that neither  $i$  nor  $j$  are themselves transport networks partnering with an interconnection platform to facilitate remote participation on that platform. Such actors do exist: IX-Reach and Atrato are two prominent instances in Europe.



an individual investment decision. For each interconnection relation  $d(i - j_k)$  in Figure 2 where  $1 \leq k \leq n$ , the contractual component of  $c(i, j_k)_{neg}$  is an independent negotiation between  $i$  and  $j_k$ . As noted in the discussion of interconnection, each decision  $d(\cdot)_{p-p}$  is not necessarily independent of the others. Networks  $i$  and  $j$  may both be clients of a transport provider  $c$ , such as IX-Reach or Atrato, that have specialized in providing, among other services, point-to-point connectivity between facilities. Given there is a single “option” with a not insignificant cost of provisioning, it is unlikely that either  $i$  or  $j$  will build to the other, ring the bell, and say, “Hi! I just built in! Would you like to interconnect?” In contrast, interconnection platforms make exploring interconnection options more tractable by inviting networks to become neighbors in a common marketplace, effectively reducing the costs of “ringing the bell.”

### 3.2 Interconnection Platforms

Interconnection platforms accrete network actors in part by advertising the diversity of actors participating in their marketplace. Interconnection platforms’ value is predicated in part on the network effects from creating a loci of interconnection: the more networks in the marketplace, the more choices an actor garners from provisioning transport to that platform. When  $i$  evaluates an established platform, a key question is: “On which platform do networks I would like to interconnect with participate?”<sup>13</sup> When  $i$  evaluates a newly developed platform, a key question is: “Will this platform attract the kind of participants I want to interconnect with, if any?” As earlier, exercising the platform option is the investment in provisioning transport to that platform. More precisely, exercising a *platform option* is the provisioning of transport to the facility hosting a platform in order to gain access to that attendant bundle of interconnection options. In Figure 2 transport is denoted by the subscript  $p - p$ . The colocation platform option comprises  $\{co_{p-p}, co_{fac}\}$  for colocation. The IX platform option comprises  $\{ix_{p-p}, ix_{fac}, ix_{xc}, ix_{part}\}$ . Note both comprise  $\{p-p, fac\}$ . These are explained in the next two sections, Sections 3.2.1 and 3.2.2, respectively. In both cases of exercising the platform option, it is assumed  $i$  has some information warrants this investment. For example,  $i$  has evaluated the bundle available at  $p$  and recognizes a subset of those that complement existing interconnection relations. These are based on  $i$ ’s objective function; such heuristics are discussed in terms of uniqueness and redundancy in Section 4.

#### 3.2.1 Colocation

Colocation facilities are commercial actors that furnish housing for routers and other networking and

server equipment. Here, colocation facilities’ services contribute to reducing the costs of interconnection by creating a common, typically carrier-neutral, facility for networks to build into in order to interconnect with one another. Professional colocation services typically include *a)* ample power, including backup; *b)* air conditioning, often including backup; *c)* various types of cross-connect cabling; *d)* space for routers and servers, typically in a hierarchy of units starting at the single rack unit, an entire rack, a cage comprising multiple racks, or suites for especially large tenants; *e)* “smart hands” labor for various simple maintenance tasks; *f)* cages and various other access control to ensure security between tenants (participants), and *g)* external security for the facility itself. Establishing residence at a colocation facility requires the purchase of some amount of space (at minimum however many units of rack space necessary for  $i$ ’s routing equipment and/or servers). Depending on the price structure and product bundles, this may include costs for power, AC, the type of cage, etc. Cross-connects are the physical links between colocation participants. These may be any number of media such as fiber or ethernet cabling. In Figure 2 of residence is denoted by the subscript  $fac$ ; cross-connects are denoted by  $xc$  and will be discussed shortly.

Assume network actor  $i$  and  $j$  are both present at colocation facility  $Q$  and they have negotiated some interconnection contract: both exercised platform option  $Q$ . Both have also completed  $co_{neg}$ . For simplicity and symmetry in this discussion  $co_{neg}$  follows establishing the platform option. This simplifying assumption will be relaxed when discussing the viability of interconnection options in Section 4. Networks  $i$  and  $j$  must now provision (presumably under the terms of the contract established in  $co_{neg}$ ) a cross connect ( $co_{xc}$ ) through  $Q$  to establish a physical link between their routers (and by proxy their networks).  $Q$  controls the colocation facility itself and typically performs all cabling between tenants.<sup>14</sup> Provisioning a cross connect between  $i$  and  $j$  completes the physical link. The remainder of the interconnection relation is exchanging routing information via some route information dissemination protocol, typically BGP ( $co_{bgp}$ ).

Negotiation, provisioning the cross-connect, and exchanging routing information  $\{co_{neg}, co_{xc}, co_{bgp}\}$  effectively exercises an instance of the interconnection option between  $i$  and  $j$  on platform  $Q$ . Costs are incurred at each point. Bargaining and measurement costs incur at  $co_{neg}$ . Provisioning of  $co_{xc}$  is considered a specific asset investment: the cross-connect consumes a router port for both  $i$  and  $j$ ; the cross-connect capacity is dedi-

13. In terms of options, selecting a subset of platforms amongst multiple platforms is itself an option. Returning to the discussion of common transport, firms like IX-Reach and Atrato offer this option.

14. In the case that, for whatever reason, a tenant’s equipment is in different parts of the facility, it may be necessary to establish a cross connect for that equipment to communicate. In the case that the tenant’s equipment is in a single enclosure, any cabling within that enclosure is typically done by the tenant, but may be outsourced to smart hands provided by the colocation facility.



cated to exchanging traffic over those two ports. Without reconfiguration and the cost of decommissioning the cross-connect, those resources are dedicated to that interconnection instance and cannot be repurposed. This will be compared with more general purpose capacity provisioning in IX interconnection provisioning ( $\{ix_{xc}, ix_{part}\}$ ) in Section 3.2.2; before making this comparison, consider cross-connect pricing structure.

Within and across colocation facilities there is variance in the cost of cross-connects. Some of that variance is due to operational and capital costs: for instance *a*) managing cables from many tenants to many other tenants in the same room, *b*) running cables from room to room, *c*) running cables across multiple floors, especially when the colocation facility is spread across multiple nonadjacent floors, *d*) between adjacent buildings, and *e*) between nonadjacent buildings in the same metro-area to name well-known sources of operational and capital costs. A number of network operators claim colocation facilities' costs are not sufficiently transparent, creating the perception that the colocation facility can charge relatively arbitrary prices. In economic terms, some colocation facilities have been accused of opaque, discriminatory rent-seeking behavior.

Opaque pricing and in general lack of knowledge about the costs of colocation facility operations creates information asymmetry problems. Consider the bargaining positions of large and small network actors. Large actors often have presence at many geographically diverse colocation facilities—these actors have a diverse platform portfolio (bundle). This gives the large network actor insight into pricing differences across differentiated colocation providers. Large actors can then leverage pricing insights to negotiate better terms on a variety of colocation services. In and of itself, this is not detrimental—large actors have greater visibility into the colocation supply market. Moreover, colocation facilities find large networks attractive because large networks, in turn, attract others that want to interconnect with these actors.

On the other hand, smaller actors face an information asymmetry problem. Absent visibility into the colocation market, smaller actors may easily pay much more than others that have better access to pricing data. Actors may have better information either through their own direct engagement with many suppliers as discussed above or through interpersonal networks like network operator groups. New market entrants may have neither the market power, market experience, or knowledge of colocation operations to bargain effectively. A strong interpretation of this scenario would claim price opacity is a barrier to comparing services as well as accessing and leveraging strategic options in the interconnection market. A common symptom of barriers to entry is price distortion. This is not a conclusive analysis: more data is needed and will be collected in ongoing work.

Provisioning a cross connect incurs either a monthly

recurring cost (MRC) or a non-recurring cost (NRC). Based on conversations with network operators, monthly recurring costs are more typical of the dominant colocation facilities based in the US; a single NRC seems more typical of Europe. Establishing each interconnection relation incurs either a single NRC or a MRC for as long as the cross connect is in service. In the case of MRC, it is assumed the service duration is the same as the life of some form of interconnection relation between *i* and *j*, given one or both parties are paying for the required cross connect.

An NRC-based cross connect seems to have more flexibility from the perspective of the platform participant. Once the one-time cost has been recouped, it may be used however *i* and *j* see fit, even if the initial cost is variable. Factors discussed above also affect the variance in NRC.

In the options framing, bilateral interconnection implemented as a cross-connect is a specific asset investment. This investment requires information sufficient to warrant this investment. For smaller actors, this investment, and by proxy investment in interconnection bundles dominated by colocation-mediated interconnection relations is confounded by cross-connect pricing variance. In contrast, IX-mediated interconnection and the corresponding capacity investment is easily repurposed (it is not as specific), has arguably lower bargaining and measurement costs, and typically has more transparent and consistent pricing structures. Taken at face value, IXes could be considered a substitute for colocation-mediated interconnection. Rather, IX-mediated interconnection is framed as a low-cost option that may be sufficient for certain classes of interconnection but also provides a means to defer contractual decisions on factors that complicate more specific asset investments. In effect, IXes are much closer to a complement than a substitute for either colocation-mediated or transit.

### 3.2.2 IXes

An IX platform is a common resource that provisions a logical path from any participant *i* to any other participant *j*. A *node* is a collection of one or more switches that themselves comprise a switching fabric, the physical component of the IX platform.<sup>15</sup> A node may be in a closet in someone's basement, in an office space with sufficient power and AC, in a rack in a colocation facility, or geographically distributed in some combination of these scenarios. The most common scenario for IXes is either a colocation facility or a facility known to provide sufficiently reliable power, AC, and outside

15. Across IX providers, a fabric housed in one location is identified by a variety of terms: node, IX, NAP (network access point), IXP (Internet eXchange Point), etc. The term node is considered the most generic and is adopted here for consistency. The point in IXP is emphasized because there is a subset of actor in the community that argue point implies a limit on the geographic distribution of switches, typically limiting these to a metro-region.

connectivity. In markets with less developed colocation markets,<sup>16</sup> they are sited in facilities that are deemed to have sufficient power, cooling, and security.

A platform often comprises multiple nodes connected via some mode of transport. Conventionally and historically, the geographic extent of a platform is the metro-region. There are interesting exceptions that vary along dimensions of demand for interconnection platforms that span regions larger than the metro, how they affect markets, and their governance regimes.

Like colocation,  $i$  must have exercised a platform option with an IX  $X$  to access interconnection options available at  $X$ . Depending on the facility,  $i$  may have to pay for hosting ( $ix_{fac}$ ). If  $X$  itself is hosted at a colocation facility,  $ix_{fac}$  is typically the same as  $co_{fac}$ , but may be less given some colocation facilities see discounts for IX participants as an investment in latent demand. In some cases, typically regions with less developed colocation markets, the facility itself is provisioned exclusively to host an IX node. In this case,  $\{ix_{fac}, ix_{xc}, ix_{part}\}$  may be rolled into  $ix_{part}$  as part of IX participation fees. Similarly, if the IX is a service of the colocation facility, the contractual elements of  $ix_{fac}$ ,  $ix_{xc}$ , and  $ix_{part}$  may be part and parcel of the same contract. Data center independent IXes are more common and thus the focus of this discussion.

IX fee structures historically comprise a membership fee and port fees. The membership fee covers participation in the IX, but not capacity. Not all IXes charge a membership fee. Port fees cover the capacity contracted on the common fabric. Typical units of volume are 100M, 1G, 10G, and, very recently, 100G. Not surprisingly, these are priced differently. For instance, the prices in Figure 1 are: Prices table! Capacity utilization is monitored by the IX. Some IXes require upgrading the port at a utilization threshold such as 80%. Others argue they trust their participants to manage their resources appropriately.

Assuming  $i$  has established itself as a participant on the IX (either as a member or a customer depending on the governance model of the IX),  $i$  establishes a cross-connect between a port on its router and a port on a switch in the IX node ( $ix_{xc}$ ). By virtue of being implemented as a common fabric, the switching fabric itself provisions the technical option for any platform participant  $i$  to interconnect with any other platform participant  $j$ . At this point,  $i$  may begin “ringing bells” to establish interconnection relations. More realistically,  $i$  already has a good idea who at the IX may be willing to interconnect and part of  $ix_{neg}$  has already taken place in some other forum. For instance, assume  $i$  participates on IX  $A$  and  $D$ , and has just completed the steps above at IX  $L$ . Further assume there is substantive overlap between the participants at  $A$  and  $L$ . It is likely that  $i$  already has contractual relations

with the intersection of  $A$  and  $L$ , so  $ix_{neg}$  on  $L$  will be an update to an existing contract rather than a new contract. This is discussed at length in Section 4. Once negotiation is out of the way, the last step (in all three provisioning modes) is to establish at BPG session, in this case  $ix_{bgp}$ . At this point the interconnection relation has been established.

For the most part, IXes have no interest in interfering with the interconnection relations that occur atop the platforms they provide. That said, IXes have experimented with various interconnection policies over the course of IX evolution. The most well-known interconnection policy imposed by the IX itself is forced multilateral peering. Under forced multilateral every participant is required to establish a peering relation with every other participant, sharing routes to all the prefixes in their administrative domain. This is typically facilitated through a route server. Part of  $ix_{part}$  would require  $i$  peer with the route server and, by proxy, with every other participant.

Forced multilateral peering is often described in Europe and the US as a tactic used early in the life of an IX to encourage reluctant participants to peer and to generate traffic on the IX. Forced multilateral is generally unpopular in the US and EU. In contrast, two of the largest IXes in South America, CABASE in Argentina and PTTMetro in Brazil, both implement forced multilateral peering in their IX ecosystems. Recall peering does not require routing traffic to the rest of the world. Forced multilateral does not preclude further negotiation between networks to sell partial or full transit on the IX.

The two most common measures of the success of an IX platform are number of members and traffic levels, either average or peak. There is some contention over the utility of these measures. As indicators, they are easily observable measures that require little normalization to compare the “success” of some set of IXes. A number of actors have indicated that, despite their widespread use, these indicators are potentially misleading. Rather, measures of outcomes are better indicators of the success of an IX. Instances of benefits not captured by either number of participants or traffic levels are: *a*) reduced latency, *b*) shorter routes to participants on the IX, and *c*) increased overall bandwidth amongst actors. These benefits are common to both IX and colocation platforms. Moreover, they are immediate. Collateral benefits that are considered to follow the development of an IX include: *a*) reduced transit costs, *b*) development of a local market for infrastructure services such as local content hosting, and *c*) further development of colocation markets.<sup>17</sup> The next section unpacks a more subtle benefit, the character of interconnection mediation in the IX and its relation to the availability and value of the options

16. Level of development should not be confounded with level of competition in that market. One can have a highly developed oligopoly, for instance.

17. See Galperin’s discussion of how IXes contribute to the telecommunications development ladder.

available.

### 3.3 Mediation and Options

Consider an IX comprising approximately 50 participant. This IX is probably not creating the traffic volume seen at large IXes with hundreds of members. Nevertheless, interviews have stress beneficial outcomes are relative. For instance, the IX may substantively reduce the costs for small participants, depending on how much of their traffic they offload to peering rather than transit. Further, participants may benefit if the IX facilitates traffic growth and better quality of experience via more direct interconnection with other participants. At the IX, once connected to the fabric, participants may interconnect with any other participant with no further mediation by the IX and without additional costs (assuming they do not exceed contracted capacity). In effect, contracted capacity and participation is general purpose: it may be used to “turn up” new interconnection relations, monitor the relative traffic volumes on these relations, and, as alluded to earlier, transition these to dedicated specific assets when warranted.

Returning to the story of our friendly peering manager ringing bells around the interconnection platform, ringing the bell implies  $i$  is ready to initiate negotiations to interconnect with  $j$ . The differences discussed above and illustrated in Figure 2 highlight how platform providers mediate interconnection option markets in terms of how links in the interconnection are provisioned and, subsequently, access to other participants. It is important to note that most interconnection platforms in and of themselves *do not* introduce new types of interconnection *contracts*. Rather, differences in access and the cost of access may yield different distributions of contract types, but has not contributed to a class of “IX-mediated contracts.”

Options are a convenient way to model the dynamic decision making processes, in particular the practice of deferring specific investment decisions until better information is available.<sup>18</sup> Figure 2 illustrates how the choice to interconnect is deferred in different interconnection mechanisms. Transport-mediated interconnection is the degenerate case: negotiation is not deferred but is necessary for any subsequent investment in interconnection. In contrast, exercising the platform option is akin to purchasing a bundle of interconnection *unexercised options* that may be exercised as information about the value of these options becomes available. Moreover, one benefit of the IX is that it reduces the transaction costs of exercising an option and potentially opens the door to better informed specific option investment. In both cases, it is important

18. There are two options perspectives one can take. Here, options are used as an appropriate modeling framework for representing dynamic decision making and investment processes that exist in the wild.

to stress that investment in the platform option itself does not establish any new interconnection relations, it simply provides access to that marketplace.

In terms of deferral, an obvious benefit of the interconnection platform is that  $i$  did not have to negotiate and pay for transport to each participant at  $Q$ . Having exercised the platform option,  $i$  is now able credibly negotiate interconnection relationships with others in the marketplace. For instance, if  $i$  is looking for transit,  $i$  may potentially negotiate better pricing based on knowledge of which other transit providers are present on the platform. Another benefit is the ability to respond to traffic growth through strategic interconnection. As traffic grows, it may be the case that  $i$  sees an increasing amount of traffic with network  $j_l$ , also a participant at facility  $Q$ . If that traffic is traversing transit for both  $i$  and  $j_l$ , at a certain threshold it will be less costly for both to interconnect (assuming the cost of the cross connect is less than the total cost of transit both are incurring to send traffic to one another). In effect,  $i$  (and  $j_l$  in this latter case) have access to the option, but deferred exercise in the form of interconnection negotiation, until both had information that warranted the transaction cost of negotiation and the cost of the cross-connect.

In the simple case of IX-based interconnection,  $i$  and  $j$  may recognize they exchange an insignificant amount of traffic and wish to shift this traffic to an interconnection platform. Relative to interconnection via colocation or IX interconnection, traffic over transit may not take the most optimal path from the perspective of  $i$  and  $j$ , almost certainly has higher latency, and has a higher probability of facing externally generated congestion.<sup>19</sup> These factors affect demand for the services. A number of interviews have described the realization of latent demand when one or more of the factors above have been removed: lower latency and less congestion have both led to an increase in traffic. The actual increase is known anecdotally and based on repeated observation and experience by some actors, but there is not a clear and easy calculation. Thus, traffic growth resulting from capacity upgrades from transit to a more specific bilateral interconnection relation *is likely*, but a clear generalizable formulation of the *magnitude is uncertain*.

Returning to interconnection provisioning, consider the simple case of exercising an option on the IX. Consider  $ix(i - j_2)$ , in which  $i$  interconnects with  $j_2$ . Assume they are confident settlement-free peering will create value for both of them. In the simple case, both have provisioned some capacity on the IX, referred to as  $cap(i)$  and  $cap(j)$ . Let  $flow(i, j)$  be the traffic from  $i$

19. Colocation- or IX-mediated interconnection is not immune to congestion, but the only sources are either  $i$  or  $j$  and on the IX, congestion elicits a signal from the IX to upgrade capacity. In contrast, in transit  $i$  and  $j$ 's traffic must content for resources with the traffic of other of  $t$ 's transit clients, creating “external” sources of congestion for their flow.

to  $j$  and  $flow(j, i)$  is the flow from  $j$ . In the case where  $flow(i, j) \ll cap(i)$  and  $flow(j, i) \ll cap(j)$  neither  $i$  nor  $j$  need to upgrade their capacity to interconnect.<sup>20</sup>

In the simple case, the only steps necessary in an IX “mediated” option is  $ix_{neg}$  and  $ix_{bgp}$ , which does not involve active IX mediation at all. By virtue of interconnecting over the IX provisioned switching fabric, the *paths* between  $i$  and any other IX participant  $j$  have already been provisioned. Exercising the option in the simple case is purely contractual ( $ix_{neg}$ ) and technical coordination between the routers at either end of the link ( $ix_{bgp}$ ). In comparison with the colocation platform, the colocation provider has been disintermediated from the interconnection *option* but not the infrastructure. In the simple case, illustrated by  $ix(i - j_1) \dots ix(i, j_l)$ , the first cross-connect established with the IX provides *access* to all interconnection options available in the IX platform and, in this example, supports  $l$  interconnection relations for no additional *platform* costs. In the case of strictly bilateral interconnection relations, the  $l$  interconnection relations will each incur transaction costs.

To elaborate, two assumptions will be introduced then relaxed in subsequent discussion. Consider the equivalent interconnection relations between  $i$  and some  $j$  in  $j_1 \dots j_l$  under colocation. The first assumption is that such a relation in the colocation platform has a significant bargaining cost, such as who should finance  $co_{xc}$ . The result is that  $co_{neg} - ix_{neg} \gg 0$ .<sup>21</sup> Second, none of the participants in the  $j_1 \dots j_l$  interconnection relations (bundle) need to upgrade their IX port ( $ix_{xc}$ ) to support the traffic in their respective relations in the IX. Under these limited conditions, disintermediating the colocation facility under the IX-mediated option has eliminated elements of the transaction costs of negotiating: the bargaining and measurement costs are completely bilateral.

Relaxing the first assumption is fairly simple:  $co_{neg} - ix_{neg} \approx 0$ . This may be the case for multiple reasons. One is mutual network development:  $i$  and  $j$  have pre-existing relations and this is part of an existing, well-known relationship with well-understood terms. This may be especially true if information revealed through an IX-mediated relation provides sufficient information to invest in a dedicated cross-connect. Another reason is the opposite end of the spectrum: both are participants in a route server and have agreed to de facto settlement-free peering with all other actors on the route server. In either case, negotiation may not be

20. The notation  $\ll$  is used because  $cap(\cdot)$  is a general purpose resource, there may be many flows on this resource. The rule of thumb for transitioning from IX capacity is if  $flow(i, \cdot) \gtrsim 30\text{--}40\%$  of  $cap(i)$ .

21. Here,  $\gg$  implies significantly greater than. In the case of settlement-free peering, it may be the case that  $ix_{neg} \approx 0$  and thus  $co_{neg} \gg 0$ . In the case of settlement-based, it would seem that if  $co_{neg} - ix_{neg} \gg 0$  and the cost of negotiating financing of the cross-connect is a significant portion of the overall transaction cost, then that negotiation may be proxy for some other element of a larger interconnection relation.

necessary. Distinguishing between the two platforms does highlight empirical questions regarding when negotiations are costly, or, framed more positively, under what criteria are bargaining and measurement costs minimized?

Relaxing the second assumption requires a subsequent investment. The platform option provides *access* to all options, but in the case of IXes, it is very likely that if all of those options were realized, a single cross-connect with the IX node would not be sufficient. Depending on capacity contracted when initially exercising the IX platform option, that may not be sufficient to support all, or even a few, of the viable interconnection options without congestion. Consider the interconnection option  $ix(i - j_{l+1})$ . Given the option will be exercised because  $i$  and  $j_{l+1}$  recognize they have a significant volume of traffic going over transit and they believe they can negotiate a better rate, it is expected  $i$  and  $j$  will also recognize that their interconnection relation will push one or both over their existing IX capacity.

As an illustrative example, assume  $i$  contracted  $cap(i)$  capacity when it exercised the IX option. Further assume  $i$  will exceed capacity contracted when exercising the platform option  $l + 1$ , i.e.,

$$\sum_{1 \dots l} flow(i, j) < cap(i) \quad (1)$$

$$\sum_{1 \dots l+1} flow(i, j) \geq cap(i) \quad (2)$$

but that  $j$  will not.<sup>22</sup> In order to complete the interconnection relation  $i$  must contract additional capacity on the IX platform, typically by either upgrading to a larger port or contracting another port. IXes enforce port capacities by shutting off participants that exceed their port capacity. In some scenarios, IXes require participants to upgrade, by either contracting a larger port or contracting an additional port. In the former,  $i$  may cancel its initial cross connect and establish a new cross connect for the larger port. In the latter case, illustrated in Figure 2,  $i$  contracts a second cross-connect from the host facility and another port at the IX.

These scenarios highlight a key benefit of IX-mediated interconnection provisioning: traffic aggregation and cross-connection utilization. Under IX-implemented interconnection options, multiple relations are multiplexed over a single cross-connect, increasing the utilization of that investment and conserving valuable router ports. Hence it being a more general investment. In contrast, if  $i$  were to establish interconnection relations with  $j_1 \dots j_l$  via colocation mediated interconnection and we assume the traffic

22. This is just a simplification.  $i$  and  $j$  needing to upgrade capacity are technically independent of one another but may confound the negotiation if one or both realize this relationship causes the other to upgrade.

levels would be the same under both modes of interconnection provisioning,<sup>23</sup> then  $i$  must provision  $l$  cross connects  $co_{xc}$  for the same traffic formerly handled by the single cross connect  $ix_{xc}$ . This is a worst case scenario, it is likely  $i$  will exchange enough traffic to warrant the capacity but absent evidence of how much, it is not clear  $i$  should invest in a cross-connect for each. Network  $i$  would also incur the cross connect negotiation costs  $co_{neg} - ix_{neg}$  for each. The prospect of this scenario is a canonical case for identifying an investment that both defers a presently uncertain investment and yields information that reduces that uncertainty, i.e., an IX-mediated interconnection relation.

This comparison and contrast does not imply one mode of implementing interconnection relations and interconnection options is universally better than the others. Under an options framework, IX-mediated interconnection is modeled in terms of the information at hand about the character of traffic exchange. Given an IX-mediated interconnection relation, participants may observe the change in traffic, contributing additional information for the next investment decision. In effect, the IX-mediated interconnection options are also an investment in information collection. Framed as investment options, staged interconnection investment reflects the dynamic decisions informed by empirical observations of traffic growth and short term growth forecasts rather than taking a long term static approach. Such a framing is consonant with interviews—ongoing work is refining these dynamics to quantify the conceptual relations developed here.

Consider the implications of the staged investment scenario. Network  $i$  and  $j_i$  have an existing IX-mediated interconnection relation. Traffic between  $i$  and  $j_i$  comprises a significant portion of their respective capacity at the IX and it is growing. Rather than upgrading their port capacity at the IX, it may be more cost effective to exercise a colocation-mediated interconnection option, provisioning a dedicated cross-connect. Staged decision-making has a number of effects: *a*) it defers capacity upgrades at the IX; *b*) it consumes an additional router port at both  $i$  and  $j_i$ ; *c*) it increases the management and monitoring costs; *d*) it frees up a portion of existing capacity contracted by  $i$  at the IX, increasing the pool of viable interconnection options with other networks on the IX platform;

23. Traffic levels between networks are not independent of the path. Rather, the higher quality the interconnection and greater the bandwidth, the greater the traffic. For instance, when a content cache is introduced in a network or IX platform, traffic comprising the content becomes greater than it was before even though no new content was nominally made available. Similarly, removing cost-based limitations on an under-utilized link will also result in a non-linear jump in traffic. These are relatively well known traffic phenomena that result in some combination of capacity and quality of service, but the precise dynamics is a dissertation in and of itself. That said, it is taken as given that traffic growth is an empirical input in interconnection decisions and benefits for platforms that allow staged investment with periods of traffic observation.

*e*) reduces the load on the IX platform itself, deferring upgrades by the IX.<sup>24</sup>

Under the scenarios described thus far, IXes are the first step in a pipeline of investment options. First stage investment provides access to a diverse set of actors for whom the value of the interconnection option is potentially uncertain.<sup>25</sup> Selecting and exercising the option contributes to an interconnection bundle in two ways: it provides the benefits of single hop interconnection and it serves to provide information for subsequent investments in that relation. The contracted capacity allows  $i$  to exercise a number of options at once, building up a knowledge base regarding with whom  $i$  exchanges traffic more than others and with whom traffic exchange is more valuable, better informing investment prioritization. Further, most member-based IXes have transparent, non-discriminatory pricing schemes, in other word, price transparency. In the case of membership-based IXes, governance norms ensure the IX provider does not favor any particular member, regardless of size or capacity provisioned.<sup>26</sup> In contrast to the variance in cross connect costs at colocation centers, IXes offer stable per port prices.

### 3.4 Options to Bundles

Thus far the discussion has focused on the mechanics of interconnection relation implementations. From the perspective of a network actor  $i$  basing decisions *only* on the information available on a single interconnection platform, with some allusions to the bundles that will be discussed shortly. More realistically, and discussed in the next section, interconnection decisions are based on  $i$ 's existing interconnection bundle (or portfolio) that may comprise multiple platforms. Further, value it places on interconnecting with actors in  $i$ 's value network. The mechanics help explain the sources of barriers to interconnection rooted in distinctions between accessibility of options and reactions to changes in traffic. From a regulation and governance perspective, it highlights loci of control and influence over entry to and degree of participation in resource provisioning dimensions of the interconnection market.

Following the pipeline metaphor above, although IX mediated and colocation mediated interconnection may

24. This is an indirect benefit for  $i$  and  $j$ , especially if, as with non-profit membership based IXes, fees are based on cost-recovery.

25. The next section will elaborate heuristics used to determine which are more viable than others based on an existing portfolio of platforms and interconnection relations.

26. In terms of infrastructure economics, this is generally a non-discriminatory pricing model. All participants are offered the same pricing structure, typically in units of ports of varying capacity (1G, 10G, 100G). Differences in participant fees are a function of capacity contracted at the non-discriminatory per unit rate. Some membership-based IXes charge an annual membership fee. This is also non-discriminatory. Prices are a product of the IXes norms regarding neutrality and is enforced in membership based IXes through monitoring by the membership (a reification of the mutuality norm) and enforced through administration by the firm (IX provider).

seem to be competing, this work hypothesizes that they are increasingly complementary. Such a hypothesis is contingent on the path dependent character of the market itself. Future work will compare the North American market (primarily the US and Canada), the EU market, and the South American market. According to some, there is, if not a monopoly on colocation mediate interconnection, a loose-knit oligopoly of colocation providers providing both colocation mediated interconnection and (commercial) IX interconnection. In contrast, the EU interconnection market, in general, has diversity in both colocation and IX platforms. Moreover, this market provides evidence of the platforms being complements rather than substitutes. Finally, the South American market is currently in the development phase after the relatively recent deregulation of the telecommunications sector.

Interviews and community observation provide some supporting evidence. In regions where IXes are predominantly data center neutral<sup>27</sup> IXes have a strong preference for facilities that can demonstrate they meet certain standards. In Europe, the diverse colocation market makes quality standards and certification a competitive advantage. In South America, this is still a preference—that said, large “national” IXes such as CABASE and PTTMetro actively encourage quality standards. Following the hypothesis above, the IX mediated interconnection market is more mature in these regions and is encouraging the development of a complementary colocation market.

Another manifestation of complementary related to transitioning IX mediated interconnection to a dedicated cross-connect is the notion of “stickiness.” As per the discussion earlier and the assumptions on cross-connect purchases, colocation facilities accrete networks that have sufficient information to warrant establishing transport to a colocation facility, hosting at that facility, and a cross-connect with *at least* one network at the colocation facility. From the perspective of  $i$ , stickiness means the capital and operational investment at a colocation facility makes leaving that facility costly. Stickiness is a function of the number of other networks  $i$  has either colocation- or IX-mediated interconnection relations with at that facility. In effect, stickiness is positively correlated with the expense necessary to transition interconnection relations to another facility. Such a move would require finding one or more facilities at which all of  $i$ 's relations at the current facility are also possible, investment in transport to that facility (assuming  $i$  is not already there), and the operations costs of reconfiguring BGP sessions and dependent network deployment decisions.

From the perspective of a network  $i$  that is only at one or a few colocation facilities, an IX makes a coloca-

tion facility even more sticky.<sup>28</sup> Having an IX present is a way for the colocation facility to accrete networks that would like to establish more sophisticated interconnection bundles than simple or n-redundant transit, but do not necessarily have sufficient information to warrant investment in cross-connects with a large number of networks at the colocation facility. Network actor  $i$  participating in IX  $A$  at colocation facility  $Q$  further compounds the stickiness of that facility. In this (fairly) common case, the colocation platform option  $co_{p-p}, co_{fac}$  is the same as  $ix_{p-p}, ix_{fac}$  in the IX platform option, making those common elements of the investment more general than simply a colocation or IX platform option. In the pipelining scenario, when  $i$  and  $j$  realize interconnection relations warrant transitioning to a cross connect,  $Q$  is then natural provider of that cross connect.<sup>29</sup> In effect, hosting an IX is a way for a colocation facility to queue up networks that are known to be in the market for diversifying their interconnection portfolio. In effect, is an investment on future latent demand. This is sufficiently appealing that a number of interviews have noted that some colocation facilities actively court the larger IXes. This ranges from providing free services to an IX as a way to entice the IX to site a node at their facility to establishing risk sharing contracts that assure the IX that the colocation facility has sufficient demand that investing in placing a node at that facility will recoup node deployment costs within a certain time period. Investments by colocation facilities that encourage an IX to site nodes in their facilities may be framed as infrastructure investment options to capture latent demand for interconnection services.

This section provided a view of interconnection largely from the perspective of a network actor  $i$ . This perspective was limited to the costs and immediate mechanics of establishing bilateral interconnection relations one-by-one as a means to focus on these mechanics and how they impact the availability of interconnection options. The next section reintroduces the notion of an interconnection bundle, highlighting the characteristics that make certain marketplaces more attractive than others based on two simple objective functions. The bundles presented in the next section comprise both strategic portfolios of *existing* relations *and* options. In particular, a key contribution of an IX ecosystem is transparency into the cost and availability of interconnection options as a means to reduce

28. Multi-node IXes may contribute to the individual value of each colocation facility but may also reduce the stickiness of that colocation facility if nodes are located at a variety of data centers, reducing the costs of moving.

29. There are certainly possible exceptions. For instance it is possible both  $i$  and  $j$  are at another facility  $R$  that offers a better cross connect rate.  $i$  and  $j$  may decide to transition the interconnection relation from IX  $A$  sited at  $Q$  to a cross connect at  $R$ . This seems unlikely, especially given that presence at more than one colocation facility implies something more than a network concentrated in one place. More likely,  $i$  and  $j$  may balance their load across some set of interconnection relations provisioned by  $A$ ,  $Q$ , and  $R$ .

27. Data center neutral means the colocation facility in which one or more of its nodes are sited does not have undue influence over its operations or strategic decision. This of course means that IX is not owned by the colocation facility.

bargaining and measurement costs.. This visibility in the option market facilitates better informed immediate (short term) investment decisions and further partitioning of relations in terms of those that may warrant subsequent specific investment decisions.

## 4 INTERCONNECTION BUNDLES

Like the previous section, interconnection bundles will be described from the perspective of network actor  $i$ . Assume  $i$  participates in some number of interconnection platforms  $p_1 \dots p_n \in P$ . Each platform  $p \in P$  is characterized by the set of interconnection options available at platform  $p$ .  $P^{30}$  will be referred to as a  $i$ 's *platform bundle*. For each  $p \in P$ , the bundle of options at  $p$  is denoted  $I(p)$ . Network  $i$ 's corresponding *interconnection bundle*  $I(P)$  comprises existing interconnection options across all  $p \in P$ :

$$I(P) = \bigcup_{p \in P} I(p) \quad (3)$$

The set  $I(P)$  represents the set of options available to  $i$  in platform bundle  $P$ , it does not distinguish which options have been exercised or not. Moreover, it is important to note that this is the set of *options*, not *instances*. For example, an option between  $i$  and  $j$  is listed once but instances may be available on any (or all) of the platforms in  $P$ . Redundancy and its implications are addressed later in this section.

In the stylized scenarios in the previous section, the discussion alluded to, but did not explicitly take into account, the relationship between an existing interconnection bundle  $I(P)$  and the potential contributions of exercising a set of platform options in  $P'$  (assuming  $P \cap P' = \emptyset$ ). When real networks consider investing in new platform options, these decisions take place in the context of *existing* interconnection relations and options. This section will focus on relations that distinguish  $P$  from  $P'$  and  $P \cup P'$ . Given an interconnection bundle comprises both existing relations and unexercised options, two notations are introduced to distinguish these.  $R_i^P$  is the set of options available in  $I(P)$  that have been exercised by  $i$ . Again, this does not indicate the level of redundancy across the platforms in  $P$ , just that on at least one of those platforms,  $i$  and  $j$  have an interconnection relation. The second notation is  $O_i^P$ , the set of options that are available, but have not yet been exercised, more formally  $I_i(P) = O_i^P \cup R_i^P$ .

The relations described here formalize some of the common comparative heuristics alluded to in interviews. Existing interviews provide a baseline, the background concepts commonly referenced in discussing the value of an IX to participants. The formalization here refines these based on information commonly

available to operators. Further, the formalization attempts to highlight different metrics that have been alluded to—future work will present these more refined heuristics to the community to understand precisely which are used in investment decisions and under what circumstances. Here, Comparative relations are used to reason about the value of exercising some set of platform options relative to an existing platform bundle and its corresponding interconnection bundle. For instance, given a platform bundle  $P$  and the opportunity to participate in platforms  $A$ ,  $B$ , and  $C$ , should  $i$  exercise the platform option with all three or some permutation? In the remainder of the paper, the platform bundle being considered, such as  $\{A, B, C\}$ , will be denoted  $P'$ .

As with any options framework, option selection depends on the objective function. Two fundamental objective functions have been articulated by a number of research subjects: uniqueness and redundancy. Uniqueness refers to how many of the options available in  $P'$  are unique to  $P'$  relative to  $P$ . In other words, which interconnection options will be available by investing in  $P'$  given an existing investment in  $P$ ? Optimizing on uniqueness may be contingent on a number of factors. Content providers and transit providers both optimize for uniqueness. For instance, a regional content provider may be interested in unique options within its existing content market. Larger (global) content providers such as Akamai, Google, or Limelight may optimize for uniqueness as a means to enter new markets.

Redundancy is also a comparative metric. Under the simplest form of redundancy,  $i$  considers how many participants in  $I(P')$  it already has an interconnection relation with in  $I(P)$ . More formally, it considers  $R_i^P \cap I(P')$ . Interconnection options may then be more precisely valued as redundancy and load balancing options. Consider the scenario in which a)  $i$  and  $j$  have a single instance of an interconnection relation in  $I(P)$  and no other instances are available in  $I(P)$ , b)  $i$  is considering investing in some permutation of platform options in  $P'$ , and c)  $j$  participates in some set of platforms in  $P'$ . If  $i$  and  $j$  are optimizing for redundancy in connectivity with  $j$ , then interconnecting on multiple platforms may be viable. If  $i$  and  $j$  do consider it valuable but do not know precisely how valuable it may be, like in the traffic growth examples from the previous section, the availability of low-cost options in the IX may provide a means to explore the benefits of redundancy. Low-cost availability allows  $i$  and  $j$  to collect better data about distributing traffic over multiple general purpose links before making the investment in a single colocation-mediated option.

The following sections refine these kinds of heuristics for reasoning about the potential benefits of supplementing an existing interconnection bundle provisioned by  $P$  with additional interconnection options provisioned in  $P'$ . Permutations of interconnection op-

30. Following the convention here, this should be  $P_i$ , but this will not be made explicit unless comparing platform bundles of two different networks, say  $i$  and  $j$ .



tions and benefits provisioned by interconnection platforms create a trade-off space conditioned on the value-proposition of the given network actor relative to the value networks they participate in. In the next section, bundles are characterized based on combinations of contracting mode and how that interconnection relation is provisioned. Refined heuristics are presented in relation to these characteristic bundles.

#### 4.1 Types of Bundles

Returning to the simple comparison in Figure 1, IX mediated connectivity is often compared to the simplicity of establishing a transit relation. The simple transit relation provides access to the rest of the world, but places all connectivity (modulo transport to the POP), routing, and upstream traffic management decisions in the hands of a single transit actor  $t$ . Simple and  $n$ -redundant transit relations represent what will be referred to as a *homogeneous transit bundle*, a bundle of interconnection relations comprising transit relations that reify a conventional multi-homing strategy. Transit bundles may be implemented using any combination of transport-mediated, colocation mediated, or IX mediated interconnection. Historically, transport-mediated and colocation-mediated are the most common. As IXes have grown, some IXes have allowed transit relations across the public fabric. In some cases IX participants have leveraged collective bargaining to negotiate better transit prices, highlighting an indirect benefit of IX-mediated interconnection.

In contrast to the contractual homogeneity of transit bundles, mixed (heterogeneous) bundles comprise multiple contracting modes: some combination of various forms of transit and peering. These can be differentiated in terms of the mix of platform types in the underlying platform bundle: what combination of transport-mediated connectivity, colocation-mediated interconnection, or IX-mediated interconnection does  $i$  use to develop its interconnection bundle? Interviews indicate a mix of colocation and IX platforms. Homogeneous platform bundles that rely exclusively on transport, colocation, or IXes are expected to be rare in the wild: rather, these bundles are ideal types that are used to establish the foundation of more realistic comparisons. They also serve to identify nuance that requires complementary empirical work through surveys and directed interviews.

A key question is what subsets of the option (commodity) space are preferred by which types of actors? A simpler, more empirically tractable question is what are the desirable characteristics of bundling strategies and what types of options contribute to those. Based on interviews, desirable characteristics are *a*) reduction in connectivity costs, *b*) latency, *c*) connectivity failure mitigation, *d*) interconnection and route potential, and *e*) potential for coordinated network planning. There are trade-offs within and across each of these.

The most conventional measure is the change in costs per volume of traffic, typically savings over the cost of transit. Many of the comparisons below take reductions in transit costs as a given. This work specifies a framework for highlighting and better understanding how IXes may or may not contribute to reducing uncertainty in transactions, thereby reducing bargaining and measurement costs. Latency has become increasingly important and is fairly well-understood. The focus of the rest of this section is to elaborate points related to interconnection and route potential as value added through flexibility and access.

#### 4.2 From Availability to Value Potential

Section 2 established the mechanics of interconnection option availability. Availability can be further refined to reflect the viability of an interconnection option, whether the option to establish redundant connectivity is available or not, the implications of different option implementations, and the implications of services such as route servers and enforced multilateral peering.

##### 4.2.1 Simple Potential

As a background concept, simple potential is the set of interconnection options available in  $P'$  but not in  $P$ . The value of  $P'$  depends on the value of:

- 1) interconnection contracts available within this bundle,
- 2) redundancy to value network complements,
- 3) traffic distribution over portfolio elements

among other items. *Interconnection potential* provides the upper bound on the number of interconnection options. Interconnection potential refers to the scenario in which every option  $o_{i,j} \in \{I(P')/I(P)\}$  is exercised. The value of interconnection potential is not necessarily the highest value bundle of relations in  $P \cup P'$ , though. It may be the case that  $i$  is looking for very specific interconnection options and the overhead of negotiating additional interconnection relations or the cost of either  $ix_{xc}$  and/or  $co_{xc}$  necessary for all possible options is greater than the value of interconnecting with those actors. Interconnection potential must account for the financial responsibilities of the contract (not all are settlement-free), the costs of platform participation, cross connect costs (if any), and negotiating costs. Interconnection potential does not necessarily result in positive value. For some actors, most notably enterprises with little interconnection experience, low traffic volume, or for whom the value density of the traffic is not sufficient to warrant the transaction costs of developing a sophisticated interconnection bundle (portfolio), anything more than simple transit may be costly.

Historically, the opportunity to engage in settlement-free interconnection (peering) with as many actors as possible has been a key motivation for joining an IX. Interconnection potential as an ideal

is closed related to two strategies used to incentivize peering: the deployment of route servers and the enforcement of multilateral peering regimes. These are discussed in the next section (4.2.2). Following a brief discussion of multilateral peering, the remainder of the section focuses on viability of bundles in terms of how they are valued.

#### 4.2.2 Multilateral Peering

Thus far, interconnection relations have been bilateral:  $i$  and  $j$  negotiate before deciding to interconnect. Multilateral interconnection mediated by a route server allows  $i$  to establish relations with every other network actor  $r$  participating on the route sever. This is typically predicated on settlement free peering.<sup>31</sup> When  $i$  peers with the route server, interconnection to all others peering with the route server is established. In this sense, the presence of a route server and participation on the route server guarantees some number of interconnection options are guaranteed relations, assuming  $i$  wishes to participate. As a result, some subset of interconnection potential may be guaranteed by the route server. Further, these are “guaranteed” options with  $\varepsilon$  bargaining and measurement costs to immediately exercise options with the entire set.

In a number of scenarios IXes have engaged in a policy of *forced* multilateral peering: any actor that participates on the IX must peer with the route server. The result is that every option on that platform is exercised. The exact reach of peering relationships varies: some require sharing only the immediate prefixes administered by a participant  $j$ , others require both those immediately administered and downstreams. In terms of interconnection potential, forced multilateral peering guarantees interconnection potential will be realized without the transaction cost of negotiation. For some large actors, forced multilateral may not generate positive value.

#### 4.2.3 Uniqueness

*Unique potential* describes precisely how many unique interconnection options  $o_{i,j}$   $i$  will have access to if  $i$  exercises platform option set  $P'$ . A simple operationalization of unique potential is denoted:

$$U_i(P'|P) = (I(P')/I_i(P)) \quad (4)$$

This is read as  $i$ 's uniqueness potential for  $P'$  given  $P$ . Formally this is the set of new options available under  $P \cup P'$ ; it may also be interpreted as the set of new networks available.  $I(P')$  is the set of options

31. Or simply peering with no regard to traffic levels, which is subtly different from settlement free. Clark et al. (2011) argue settlement free is predicated on approximately equal such that the value of the difference in traffic levels is lower than the operational costs and transaction costs of determining the actual difference. Peering over the route server relaxes settlement free to include situations where monitoring traffic from  $i$  to  $j$  is not worth the operations and transaction costs regardless of the level.

available via the new bundle minus<sup>32</sup> those that are available in the existing bundle  $P$ . Note that the right-hand-side is the set of options *available*, not just those that are exercised in  $P$ . For example, if  $o_{i,j} \in O_i^P$  (an unexercised option in  $P$ ), it *will not* be in  $U_i(P'|P)$ . The set represents opportunities unique to  $P'$ . Following Equation 4,  $|U_i(P')|$  represents the total number of new interconnection options that would be made available, not instances. It does not consider the topology or possible redundancy.

Consider

$$N_i(P'|P) = \sum_{p \in P} |I(P')/p| \quad (5)$$

In Equation 5, for each platform  $p \in P$  the term  $|I(P')/p|$  is the number of new options available in  $P'$  relative to  $p$ . In contrast to  $|U_i(P')|$ , summing these individual terms intentionally “double counts” instances to highlight potential redundancy amongst those unique options. Equation 5 still doesn't convey topological information, but may be useful for  $i$ 's coarse-grain prioritization.

A simple *aggregate* representation of the redundancy of paths to unique networks in this bundle is:

$$RU_i(P'|P) = \frac{N_i(P'|P)}{U_i(P'|P)} \quad (6)$$

Such a ratio gives a hint at the redundancy intrinsic in the unique set of networks accessible in this bundle. The higher  $RU_i(P'|P)$ , the greater the potential redundancy available in  $P'$  to those options that are unique to  $P'$  relative to  $P$ . Redundancy will be discussed in more general terms shortly.

Assuming all potential interconnection relations have the same value, a simple indicator for unique potential is  $|U_i(P'|P)|$ . All interconnection relations do not have the same value, though. To reason about this differentiation this section will assume a simple vector of weights  $K_i^P(P')$ <sup>33</sup> whose indexed values  $k_1 \dots k_{|U_i(P'|P)|}$  correspond to the value of options  $o_{i,j_1} \dots o_{i,j_{|U_i(P'|P)|}} \in X$ . Following the discussion of interconnection potential, if every potential interconnection option were assumed to have a positive  $k$ , Equation 7 would provide a starting point for understanding how differentially valued interconnection options affect the value of a potential bundle.

$$V_{pot}(U_i(P'|P)) = \sum_{k \in K_i^P(P')} k \quad (7)$$

Rather, again following the discussion of interconnection potential, some  $k$  may have negative values.

32. Here “/” denotes set subtraction, i.e. the left set except those in the right set.

33. Note that  $K$  is parameterized by  $i$ ,  $P$ , and  $P'$ .  $K$  is a placeholder. That said, this work assumes that the decisions regarding connectivity between network actors are driven by their value proposition and the selection amongst potential bundles is a projection of their risk profile.

Instead, Equation 8

$$V_{ben}(U_i(P'|P)) = \sum_{k \in K_i^P(P') | k \gg 0} k \quad (8)$$

is more likely. Note that the summation is now conditioned on the value  $k$  of the interconnection option being substantially greater than 0 ( $k \gg 0$ ). In effect,  $k$  is, unsurprisingly, a filter on which actors  $i$  finds it beneficial to interconnect with and thus with whom  $i$  will find it valuable to engage in negotiations over exercising an option with. As implied by the uncertainty in interconnection described in the previous section, it may be obvious that elements of  $K$  are known with varying degrees of certainty.

Hypotheses around the bounds of  $k$  and the threshold of  $\gg 0$  can lead to further insights. Understanding how  $k$  is valued given different bundles and at different stages of an interconnection relation between  $i$  and  $j$  will ultimately be better informed by directed empirical queries in future work. That said, the relations  $V_{pot}$  and  $V_{ben}$  can provide useful conceptual comparators. For instance,  $V_{pot} = 0$  could result from all the relations being settlement free or the value of various relations canceling one another out. An interesting, albeit difficult to obtain, metric would be the distribution of  $V_{pot}$  across IX participants.

Comparing bundles requires not only comparing  $P'$ , but subsets of  $P'$  and identifying the contribution of each platform  $p \in P'$  to uniqueness and redundancy. Parameterizing with a single platform is akin to evaluating the marginal contribution of that IX. The entire trade-off (commodity) space would be represented by the powerset of  $P'$ . Treating this as a static optimization problem, if  $K$  were available for a given bundle it may be possible to find the subset of  $P'$  that has the greatest  $V_{ben}$ . This may give some insight into the investments that may be promising but exercising interconnection options is not guaranteed. Realistically, exercising more than one or two platform options at once<sup>34</sup> may not be feasible (for financial and bureaucratic reasons) for any but the largest network actors.

In addition to limitations on exercising platform options, in terms of the framework presented here, treating the powerset of  $P'$  as a space over which a single optimization may be invoked is a static analysis that ignores the dynamic character of interconnection investment and the benefits of incremental investment. As indicated in the discussion in this section and the previous, the benefits of an option framework is that it highlights information deficiencies and provides a framing that stands those deficiencies up as first class decision variables (exercise the option now or later). As per the fundamental notion of an option, the key benefit of is the ability to defer decision until more complete information about the investment is available. In this

context, option investment creates a feedback loop: investment in a platform provides access to options, the exercise of which will provide further information regarding whether one should further invest in IX capacity or whether one should invest in cross-connect capacity. Framed as options within a feedback loop, a more complicated, but more realistic approach is to evaluate sequential option investments that balance stable connectivity with the data collection necessary to inform future investments. Framed this way, the objective is to understand the variety of critical paths through the powerset of  $P'$ , updating priors along the way. In contrast, framing the optimization problem as sequence of option investments highlights the role of asset specificity in the process of developing interconnection bundles.

#### 4.2.4 Redundancy

Redundant connectivity to a set of networks is the second objective function described in interviews. *Redundancy potential* is the number of instances of interconnection options in  $P'$  that are redundant with existing instances of interconnection relations in  $P$ . Consider  $i$  has interconnection relations with  $j \dots j_m$  in  $P$  and  $j_1 \dots j_m$  participate in some number of platforms in  $P$ . In this case,  $i$  has a redundancy potential of at least  $m$ , depending on how many instances of each unique interconnection option are available in  $P'$ . In this case, an interconnection option is refined into a redundancy option: the opportunity to invest in connectivity resilience and/or the opportunity to distribute traffic across multiple platforms. While resilience and traffic distribution are two sides of the same coin, they will be considered separately for conceptual clarity.

In the case of the resilience, redundancy options provide a much finer-grained set of design options than n-redundant transit. Consider a mixed bundle of interconnection relations comprising transit complemented by peering relations over an interconnection platform. The baseline (*simple heterogeneous*) case can be stylized from the original impetus of the IX: keeping geographically local traffic local. Consider a IX  $x$  and a network actor  $i$  that participates in  $x$ . For this stylized example, assume the following:

- 1)  $i$  has a set of interconnection relations with some set of networks  $J = \{j_1 \dots j_m\}$
- 2)  $i$  has a transport relation  $c_t$  which provides connectivity for  $i$  to a POP of its simple transit relation  $t$
- 3)  $i$  has a transport relation  $c_x$  which provides connectivity for  $i$  to IX  $x$
- 4)  $c_t$  and  $c_x$  are independent, and thus fail independently

This simple bundle is susceptible to:

- 1) Failure of  $c_t$ .  $i$  remains connected to  $J$  via  $x$  but loses connectivity to networks over the transit relation (rest of the world minus J).

34. In financial terms, in a single period.

- 2) Failure in  $t$ .  $i$  is disconnected from all but  $J$  via  $x$ . In a more realistic case, failure in  $t$  may be partial, disconnecting  $i$  from only a subset of the networks not available at  $x$ . In either case,  $i$  relies on  $t$  to repair the failure.
- 3) Failure of  $c_x$ .  $i$ 's connectivity is not affected since  $t$ , as transit, should provide connectivity to  $J$ . That said,  $i$  is now paying transit prices for traffic to and from  $J$ .
- 4) Failure of  $x$  has the same effect as failure in  $c_x$ — $x$  must pay transit prices to  $J$  and are now susceptible to failures in  $c_t$  and/or  $t$ .

As IXes became a more common interconnection platform, participants shifted more of their traffic from transit to the IX platform. Interviews have referenced situations where new participants have shifted a wide range of their traffic to the IX. In a case study on E4A, an Italian ISP, published by the LINX, 90% of E4A's traffic traverses the exchange.<sup>35</sup> The E4A case also highlights the value of engaging at multiple exchanges: according to that case, E4A participates in 28 IXes. E4A illustrates a much more sophisticated instance of the simple redundancy instance above. Rather than having redundancy in  $J$  at a single IX, E4A has some degree of redundancy across many IXes around Europe. Taken together with the volume of traffic over IX platforms, interconnection platform management places control of redundancy and traffic paths under the control of the network actor, not the transit provider.

One interpretation of interconnection platform contracting versus transit is what portion of your service production an actor is willing to outsource. Internet packet delivery in general is a game of mutual reliance and best effort deliver. Interconnection is a topologically localized version of that game. When a network actor engages in a transit relation, that actor outsources management of network failure, redundancy, path efficiency and selection, and congestion management to the transit provider. Strategic interconnection platform management and the resulting interconnection bundles are a means to selectively in-source decisions related to networks critical to participant's value proposition. This is evident for conventional ISPs and infrastructure providers, and anecdotally true for more narrowly scoped network actors.

Two types of redundancy can be identified in a connectivity bundle: *unique redundancy potential* and *existing redundancy potential*. Both are background concepts that can be refined into specific indicators. One aggregate operationalization of unique redundancy potential was addressed in the definition of  $R_i^P$ , the ratio of the number of instances not in  $P$  to unique options. The following indicators further refine the

notion of redundancy. Aggregate and option-specific indicators are specified below.

*Aggregate unique redundancy potential* conveys the total potential instances of redundancy options that exist amongst networks unique to the bundle being evaluated  $P'$ . Unpacking this definition, consider the following assumptions:

- 1) there are  $n$  options unique to  $P'$  (not available in  $P$ ),
- 2) amongst those, there are  $m > n$  instances available across the platforms in  $P'$

As such, there is some redundancy within that set of options unique to  $P'$ .<sup>36</sup> Before jumping into option-specific notions of redundancy, existing redundancy potential is defined.

*Existing redundancy potential* conveys how many instances of redundancy options exist in  $P'$  that have been exercised in  $P$ . In terms of a bundling strategy, this measure of  $P'$  evaluates how it contributes to the resilience of existing relations and/or how  $P'$  could contribute to mutual network planning such as load balancing or further localizing traffic through route selection across platforms. *Aggregate existing redundancy potential* conveys the total potential instances of redundancy options that exist amongst options that have been exercised in  $P$  and that are available in  $P'$ . Using the notation defined earlier, the set of redundant options is  $E = R_i^P \cap I(P')$ . The aggregate number of redundant options is the total number of instances of options across  $E$ .

A measure that can help understand redundancy is how much a network has diversified their interconnection bundle across interconnection platforms, in particular, across IXes. Given a bundle of IX platforms, a network  $j$ 's diversity score is the number of those platforms that  $j$  participates in. The term diversity is used because it represents the diversity of each network  $j$ 's portfolio of IXes with respect to a given bundle. The diversity score is denoted  $d(j|P)$ , which is read "the diversity score of  $j$  given platform bundle  $P$ ." For instance, consider network actor  $i$  and its existing platform bundle  $P$ , and some network  $j$  that participates in a subset of the platforms in  $P$ . For  $i$ ,  $d(i|P)$  is by definition  $|P|$ . For  $j$ ,  $d(j|P) \leq |P|$ . The function  $inst(o_{i,j}|P)$  is the number of instances of the option between  $i$  and  $j$  exist in  $P$ . Further,  $1 < inst(o_{i,j}|P) \leq d(j|P)$ , indicating that there may be a mix of exercised and unexercised instances of that option  $o_{i,j}$  in the platform bundle  $P$ . In terms of asset specificity, a relatively high diversity score for a network is considered an indicator of reducing asset specificity and an investment in flexibility (the ability to repurpose resources).

36. Note if  $m = n$ , there is no redundancy. Further, it is a contradiction for  $m < n$ : this would imply that for some  $m-n$  options supposedly available in  $P'$  there are no instances of the option that can be exercised. For an option to exist, at least one instance of the option must be available in a platform bundle.

35. In future interviews and surveys, IX members will be asked what proportion of their traffic traverses various interconnection platforms.

The diversity score provides a mechanism for specifying the option-specific operationalizations of unique redundancy potential and existing redundancy potential. Option-specific operationalizations indicate the uniqueness and redundancy of particular options within a given bundle, opening the door to understanding marginal contributions of options and relations to bundles. The difference between unique and existing is whether diversity is conditioned on  $P$  or  $P \cup P'$ . Given the option in question is  $o_{i,j}$  and that  $o_{i,j} \notin I(P)$ , the *option-specific unique redundancy potential* is defined as simply the diversity score of  $j$  given  $P'$ ,  $d(j|P')$ . For existing redundancy potential, given the option in question is  $o_{i,j}$  and that  $o_{i,j} \in R_i^P$ , the *option specific existing redundancy potential* is the diversity score of  $j$  given  $P \cup P'$ ,  $d(j|I(P) \cup P')$ . In these cases the indicators are very literally *potentials*, they reflect the total existing instances of unexercised options and exercised options (existing relations). In the next section, this notion is refined to differentiate between those that include exercised options as an indicator of the viability of exercising subsequent instances of the option on additional platforms.

#### 4.2.5 Moving from Potential to Realized Interconnection.

Potential interconnection is the number of possible interconnection agreements given a connectivity bundle. The *realized* interconnection is the number of those potential agreements that are actual realized by  $i$ . The discussion of  $V_{pot}$  and  $V_{ben}$  highlights that the ideal set of options that one may wish to exercise is not necessarily the potential upper bound. When considering a platform bundle, a natural selection process is to consider those platforms where networks that  $i$  already has interconnection relations with participate. In effect, within a potential bundle  $P'$ , those platforms  $p \in P'$  with the highest marginal existing redundancy potential may be a useful indicator of the viability of options in that platform.

Within the set of options referred to by  $d(j|P \cup P')$ , the greater the number of relations, the greater the viability of the unexercised options. Moreover, one already has mutual information that can potentially reduce bargaining costs. This notion of viability is based on publicly available information from data sets and the definition of existing ( $o_{i,j} \in R_i^P$ )—the more options have been exercised, the more times  $i$  and  $j$  have engaged in  $ix_{neg}$ , the more information they have about their traffic exchange, and the more likely they will recognize a benefit to subsequent interconnection agreements.<sup>37</sup> As a first cut, comparing the marginal (aggregate) existing redundancy potential can highlight which platform provides the most options  $i$  already has a relationship with.

37. There are certainly exceptions. For instance, if a particular interconnection agreement would make a peering agreement unacceptably asymmetric,  $i$  and  $j$  may not choose to exercise that option.

Simply selecting the platform with the greatest marginal existing redundancy is not always the best decision. For instance, while there may be more redundancy in the sense of the number of options, that combination may not be the set of options that are most valuable to  $i$ . The marginal analysis is perhaps better informed by using  $K$  as a filter. It may then be the case that filtering based on  $K$  highlights that the marginal existing redundancy in another platform has a higher value even though it has nominally fewer redundancy options. Another possibility is that some combination of platforms may yield redundancy options that are more valuable to  $i$  than the single platform with the largest marginal redundancy. Identifying and exploring the distribution of marginal values is part of ongoing work.

These scenarios can be identify by unpacking the aggregate measures for a platform or bundle of platforms and evaluating the option-specific indicators. The option-specific indicators highlight precisely which relations are available via what platforms. For example, consider the following situations

- 1)  $i$  has existing interconnection relations with  $j_1 \dots j_n \in P$
- 2)  $i$  wishes to exercise additional redundancy options with  $\{j_1 \dots j_n\}$
- 3)  $i$  is considering a potential platform bundle  $P' = \{A, B, C\}$
- 4)  $j_1 \dots j_{n-2}$  have corresponding  $K$  values of  $k_1 \dots k_{n-2} = 1$
- 5)  $j_{n-1}$  and  $j_n$  have  $k$  values of 20 each

By virtue of having existing relations with  $j_1 \dots j_n$ , it may be more like that  $i$  can establish subsequent instances of interconnection via other platforms. In this example, the weights on  $j_{n-1}$  and  $j_n$  will clearly skew the value of a platform bundle. If these two participate in all of the platforms on  $P'$ , it is simply a matter of selecting the platform(s) with the greatest aggregate existing redundancy. If on the other hand,  $j_{n-1}$  participates in  $A$  and  $j_n$  participates in  $C$ , with the remaining  $n - 2$  participating in  $B$ , it is clear that the aggregate expected redundancy is potentially misleading. It becomes more misleading if the  $n - 2$  are distributed such that some participate only at  $A$ ,  $B$ , or  $C$  but  $B$  remains the highest aggregate expected redundancy.

#### 4.2.6 Routes and Existing Redundancy

In a number of interviews, IX participants and IX operators have argued that participation on the IX provides participants with visibility into a wide variety of routes to different networks. In the case of simple transit,  $i$  simply “points default” at the transit provider and trust the transit provider to make decisions amongst potentially different routes to the same network destination. Network  $i$  may be paying  $t$  for transit, but so is some large number of other actors. Transit provider

$t$  must optimize for its own route selection metrics, which may not be aligned with the choices  $i$  (or any other of its customers) may make. In effect,  $i$  has outsourced its routing decisions to  $t$ , or in the case of  $n$ -redundant transit, some set of transit providers. Interviews imply that interconnection bundles in the IX provide greater control of route selection. In effect, managing a diverse set of interconnection relations gives one greater visibility into the routes a packet will likely take and, given a set diversified along the lines of redundancy, allows actors to select amongst different routes.

Consider the process of route dissemination from the perspective of  $j$ . When  $j$  participates in platforms that do not force multilateral peering, it can be very selective about the routes it advertises to its networks and its downstream networks. For instance, at platform  $p_1$  network  $j$  may advertise route  $j - k - l$  to network  $k$  and on platform  $p_2$  it may advertise route  $j - m - n - l$ . In this case,  $j$  may be making an internal decision regarding the costs or potential revenues that may be garnered from these two routes and the actors they are advertising those routes to in that geographic region.

From the perspective of network actor  $i$ , the platform  $P' = \{p_1, p_2\}$  provides redundant instances of options with  $j$ , through each of which  $l$  is reachable. From  $i$  perspective, if it has exercised the option with  $l$ , it now has two different paths to select from.<sup>38</sup> Under one scenario, both  $i$  and  $j$  have redundant interconnection relations with one another, protecting both against a failure of one of them. Further, from the perspective of  $i$ , it may choose to send all of its traffic over one or the other path. An option that has been offered by a number of actors is that rather than  $i$  and  $j$  acting opportunistically,  $i$  and  $j$  actively coordinate the traffic sent over the two common interconnection relations. In effect, redundancy not only protects against failures, it also opens the door for coordinated network planning.

Earlier, the notion of general versus specific asset investment was discussed as a means of differentiating between IX-mediated and colocation-mediated interconnection options. In the case of coordinated network planning, the IX provides a low-cost mechanism for exploring these benefits. In terms of information about a potential interconnection partner, in some cases  $i$  may know little about  $j$ 's credibility or whether  $j$  will adhere to the terms of traffic exchange agreements. These agreements incur some cost to monitor and enforce. By exchanging traffic over the general asset, IX provisioned options,  $i$  and  $j$  can develop a sense of the others' behavior as an interconnection partner and, more precisely, a partner in mutual network planning. Depending on the cost thresholds, it may not be worth it for  $i$  and  $j$  to invest in colocation-based

interconnection. Again, IX connectivity has arguably lowered negotiation costs ( $ix_{neg} < co_{neg}$ ) and allowed deferral of investment in a specific asset until further information is collected on the relation.

## 5 CONCLUSIONS AND ONGOING WORK

Comparing IX-based connectivity with transit purely on a price comparison is a quick and convenient indicator, but does not capture the more nuanced benefits of IX participation. Interviews have referenced two common strategies when including IXes in the mix of platforms for strategic interconnection bundle development: the value of unique networks and the value of redundancy. Distilling these and the differences in mechanisms across modes of interconnection, the model presented here highlights the dynamic character of the interconnection market from the perspective of interconnection mechanisms and the underlying platforms that provision those mechanisms.

IXes play a distinct role in providing general purpose interconnection mechanisms. Among other things, these mechanisms and the options engendered facilitate transitioning from functional, yet simple, interconnection bundles such as homogeneous transit bundles to more sophisticated bundles comprising a variety of relations across multiple types of interconnection platforms. An options framing highlights a key insight into deferral of specific investment facilitated by the IX ecosystem: IX-mediated interconnection can be leveraged as a low-cost general asset investment that allows the deferral of potentially more expensive specific asset investment. A complementary transaction cost framing provides the basis for developing hypotheses around how different interconnection modes and the resultant strategic bundles affect bargaining and measurement costs. Ongoing work will attempt to elicit more concrete valuations and parameterizations of these costs.

On the face of it, IX interconnection is little different from the relations available at colocation facilities. A key difference is precisely how interconnection is provisioned: how specific an investment is necessary, what information is available, and, most importantly, whether networks have options to defer decisions under conditions of uncertainty. Taken together, actors previously limited to transit bundles have the opportunity develop strategic interconnection bundles that not only reduce connectivity costs, but also provide the opportunity to reduce bargaining and measurement costs by experimenting with interconnection over the general purpose (public) resource before making the more specific asset investment in the (private) resource. The similarities between colocation and IX mediated interconnection—low-latency, dedicated capacity with only a single hop between networks  $i$  and  $j$ —allow actors to evaluate the benefits of interconnection via the general resource (often referred to as the

38. If both paths were seen by the same router, it would select amongst the two paths. Based on the BGP protocol, the rule is to select the shorter path. That is not the case here, though. Rather,  $i$ 's router at  $p_1$  sees  $jkl$  and  $i$ 's router at  $p_2$  sees  $jmn$ .

public platform) without the specific asset investment in a dedicated port and cross-connect. Rather than a simple story of one mode of connectivity substituting for another, the interconnection market is better represented as a story of selecting amongst multiple modes of connectivity that complement different states of interconnection bundle development. One hypothesis being further refined in ongoing work is that IXes contribute to learning effects in a market, that the public interconnection fabric is a venue for developing strategic bundling skills via a general purpose resource, improving the efficiency of the interconnection market as a whole. It is thus not surprising that incumbent carriers are actively wary of platform that may well facilitate learning effects for challenger sets.

IX-mediated interconnection options are key to deferring specific investment in the framework presented here. The IX platform provides a low-cost, reusable resource for exploring interconnection relations with a wide variety of other participants. Current interviews provide preliminary (but limited and generally anecdotal) support for the hypothesis that IX-mediated interconnection is low-cost in terms of bargaining and measurement:  $i$  and  $j$  do not need to negotiate payment for the cross connect and if the interconnection relation does not work out, the capacity can be repurposed with little additional cost. Further, the general purpose resource uses fewer router ports. For both small actors that do not have the operations capacity and larger actors looking to economize operations management and ports as a resource bundle, this is a significant benefit.

IX-mediated interconnection is not a universal substitute for direct interconnect via a colocation-based cross-connect. In some cases, it is a gateway to explore more direct options, in other cases IX-mediated interconnection is sufficient—ongoing work is building on the distinction to develop a survey of existing and potential IX participants to better refine the distinction and how it manifests in practice. IX-mediated interconnection allows potential interconnection partners to explore indicators of traffic patterns, most notably growth, in an environment that has the same characteristics as direct interconnection. In contrast to transit, which is optimized from the perspective of the transit provider's value-proposition in relation to its large number customers, both the IX and colocation facilitate single-hop interconnection relations.<sup>39</sup> As an "experiment" the IX as a general purpose resource has the same latency and traffic growth potential as a direct connection (modulo limits on IX contracted capacity, discussed earlier) and provides an opportunity to observe whether traffic does in fact grow as expected. Making the immediate transition from transit to a dedicated cross-connect has varying degrees of

uncertainty that creates measurement costs. Running the experiment eliminates those costs in lieu of the opportunity to sample actual traffic.

In terms of bargaining, one hypothesis is that this eliminates some of the need to condition contracts on traffic growth or on ratios based on traffic exchange over transit. Rather, the contract may be conditioned on gathering better information to inform subsequent decisions.<sup>40</sup> Such conditioning is hypothesized to reduce measurement costs and contributes information to both sides of the connection that can be leveraged in subsequent bargaining if additional investment is warranted. If growth warrants investment in a more specific asset, here colocation or even transport-mediated interconnection, both actors now have better information going into the bargaining process. Moreover, both actors have some additional degree of experience dealing with the other, contributing again to learning effects.

Thus far, the benefits here tell a rather pretty story of mutual network planning. An empirical question for ongoing work challenges the assumptions of mutual development many of the hypotheses and specifications presented here take for granted. Does additional experience dealing with a potential interconnection partner always engender lower subsequent bargaining costs? Under what conditions does experience engender *higher* bargaining costs? A number of the interviews have implied lower costs, but there are dissenting voices. Ongoing work will delve into the factors and contexts that limit the potential for mutually network planning.

For instance, are their scenarios where competitors may benefit from mutual planning? Consider a hypothetical set of hosting providers and their clients. Although hosting providers nominally compete, many hosting customers diversify for redundancy. Hosting providers themselves may recognize some of their clients want diversity, thus making it beneficial for hosting providers to interconnect rather than pay transit or pass the costs of transit along to their customers. In contrast, it does not seem likely CDNs would have similar incentives. As evidenced by Level3-Cogent and other cases, lower bargaining costs is not always the case, even when the actors have longstanding relationships. In one developing region, a number of rural networks have refused to interconnect with one another across the IX because they perceive interconnection conferring advantage onto the other. Both pay high transit costs and both exchange traffic with one another via transit.

Another hypothesis related to mutual planning is

40. As per discussion in fieldwork and a survey by Woodcock, many interconnection agreements are based on handshake, not a formal agreement. In this work a contract refers to either an informal or formal contract. Although there are certainly differences between the two, the focus here is on how opportunities for development of strategic bundles affect the *terms* of the contract, not the manifestation or the formality of the contract or its enforcement.

39. Single hop refers to a single hop between  $i$  and  $j$ , not between the downstreams of  $i$  and or  $j$ .



the notion of open book coordination from supply network strategy. For instance, some CDNs have offered to share information about cache management in order to facilitate coordinating traffic with access networks. On the face of it, this can be beneficial for both parties. Lightweight IX-mediated interconnection may be one vector for developing these mutually beneficial relations, especially given IXes are a very favorable location for CDN caches. In supply network strategy, this is akin to an “open book” contracting model where production and logistics constraints are shared amongst actors in a supply chain to facilitate better planning and to avoid the effects of poor signaling, such as the bullwhip effect. One hypothesis is that general purpose, lightweight interconnection relations facilitated by IXes not only facilitate reducing bargaining costs and measurement costs, but they can also facilitate better communication and mutual planning. Existing interviews allude to such arrangements; ongoing work is developing these anecdotes into directed interviews and survey questions to elicit a sample of instances. Such a sample is expected to provide more substantive evidence for (or against) this form of coordination and the conditions under which it does and does not occur.

The plight of rural networks in developing regions provide a useful segue to hypothesis about the value of IX interconnection and market saturation. One explanation of the case above, where two rural networks face high transit costs to exchange traffic, is that both face tight margins and thus perceive any advantage conferred to the other as untenable. In many other scenarios, small actors in developing regions are happy to band together to get out from under the high costs imposed by a common transit provider (this is the case in Argentina, Kenya, and a number of other developing regions adopting IXes as a catalyst for growth). One hypothesis that emerged in private conversations was the willingness to engage in IX development and subsequently IX-mediated interconnection as a function of market saturation: historically networks have been happy to enter into cooperative agreements when there is some combination of *a*) a relatively uncrowded market; *b*) untapped markets (green fields) nearby, creating the perception of a surfeit of untapped demand; *c*) high transit prices that incentivize participation in the IX. Under a saturation hypothesis, one would expect *a*) development of and growth in participation in IXes and colocation in developing regions<sup>41</sup>; *b*) leveling of participation by local actors at mature IXes in saturated markets; *c*) investigation of less saturated

markets by both new and “incumbent” IXes *d*) efforts to identify latent demand for IX connectivity in mature markets that do not have IX coverage. In the cases where one expects growth, a general assumption is that these actors are “in-sourcing” critical relations back from transit given the opportunity (option). In developing regions, a concurrent phenomena has been a drop in transit prices as IXes are introduced. Interviews have alluded to large drops in transit prices where IXes have deployed into incumbent carrier territory; instances in Argentina have been confirmed by ? (?).

Points b-d can be viewed as saturation of existing IX markets and identification of new markets, especially by existing IX providers. These hypotheses are interesting in themselves and will be used to help frame the value of IX interconnection options under each of these general sets of market conditions. For instance, one characterization of the US interconnection market is that interconnection options are largely concentrated in a relatively small number of colocation facilities in a few large metro-areas: the market may be characterized by a (relatively) small number of high-density colocation platforms. The EU is characterized as having a larger and more diverse (in terms of platform providers) set of colocation providers accompanied by a diverse set of colocation-neutral (and independent) IXes: the market has a larger number of lower-density interconnection platforms that offer a mix of IX- and colocation-mediated interconnection. A number of developing markets have yet to go down either path: it is an open question how those markets will develop interconnection infrastructure, but there are lessons to be garnered from both the US and EU markets and the implications of how interconnection option provisioning affects growth and diversity.

As noted in the abstract and introduction, the objective of this work is to develop a working specification that helps refine qualitative evidence into relations that provide insight into how to further operationalize the value of IX participation. The hypotheses offered in the body of this work and this concluding section are a product of that effort. The options framework and mechanics of different modes of interconnection option provisioning provide the basic processes that are hypothesized to be key variables in the interconnection market writ large. In particular for one of the motivations, interconnection market growth in the developing world, the implications for the learning effects on strategic interconnection bundle development are key. As may be obvious and noted throughout, this specification is *one* possible realization of background concepts elicited in interviews—the specification will be subject to validation in ongoing work along with the empirical questions it has inspired. For instance, the previous section moved from aggregate measures of uniqueness and redundancy to a fine-grain notion of marginal utility of a given interconnection option

41. There are many other factors that affect IX development in a region, one of which is market and political power of the incumbent, operational capacity of existing actors in the market, existing of a third party political entrepreneur with sufficient credibility to establish trust amongst nominal competitors, and stable infrastructure such as transport and power. This hypothesis is necessarily simplified: ongoing work has identified a number of cases that will explore the confluence of these factors with IX development and growth.

relative to a bundle. While it is clear this kind of analysis is possible, it is unclear how widespread this level of detailed analysis is in the interconnection market.

A number of actors have alluded to their own models they use to value interconnection relations. Beyond describing the general heuristics and particularistic instances, these actors considered their model to be a source of competitive advantage and did not share detail sufficient to reverse engineer those models. These types of models are not surprising in and of themselves, but a key question for this framework is what detail of analysis is sufficient for developing strategic bundles? Cost data may provide some bounds: in some cases it is expected to be clear that a given bundle is sufficiently valuable,  $V_{ben} \gg 0$ . Data collection and analysis in ongoing work will try to characterize the scenarios (based on actor value proposition and objective function) under which different granularities of analysis are useful. For instance, under what conditions is  $V_{ben} > 0$ ?  $V_{ben} \gg 0$ ?  $V_{ben} \gg\gg 0$ ? As noted earlier, precisely what constitutes the difference between  $>$ ,  $\gg$  and  $\gg\gg$  is an empirical question. Further refinement of the specification in the earlier sections could be used to develop a full-factorial simulation that identifies spaces where  $V_{ben}$  crosses some threshold—such a model could even be parameterized by publicly available data on pricing such as in Figure 1. Such an analysis would not tell us which of those “feasible regions” exist in the wild, in what proportion to others, their significance relative to others. Moreover, such a model would be speculation until further validated.

The next step in this work will be to present the refinement of these heuristics to IX operators and participants as the first step to eliciting validation of the relations specified here. This will be a combination of conference presentations, interviews with established research subjects, and surveys. Transit, transport, and IX cost data will also be integrated to develop a “baseline” notion of what constitutes the value of IX connectivity relative to colocation and transit. Interviews will attempt to identify and validate archetypal strategic bundles that have been discussed in private conversations.

While there are a number of empirical questions remaining, the contribution of this specification is to highlight the dynamic character of strategic bundle development. In particular, although superficially a simple difference in interconnection mechanism, the opportunities provided by IX-mediated interconnection options are argued to have potentially substantive benefits in terms of bargaining costs, supplanting measurement costs with an empirical sample of traffic behavior, and a broader opportunities for learning effects. Taken together, these arguably have the potential to reduce some of the information asymmetries in the interconnection market, resulting in a more efficient and competitive market. Ongoing work will refine and

validate these specifications and lines of reasoning to more precisely delineate the scope, degree and most ambitiously, the magnitude, of the effects of IX-mediated interconnection in their respective markets.

## REFERENCES

- Adcock, R., & Collier, D. (2001, September). Measurement Validity: A Shared Standard for Qualitative and Quantitative Research. *The American Political Science Review*, 95(3), 529–546. Retrieved from <http://www.jstor.org.libproxy.mit.edu/stable/3118231>
- Clark, D., Lehr, W., & Bauer, S. (2011, 23–25 September). Interconnection in the Internet: The Policy Challenge. In *Proceedings of the 39th Research Conference on Communication, Information and Internet Policy*. George Mason University, Arlington, VA.
- Faratin, P., Clark, D., Gilmore, P., Bauer, S., Berger, A., & Lehr, W. (2007). Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy. In *Proceedings of the 35th Research Conference on Communication, Information and Internet Policy*.
- Jensen, M. (2009). *Promoting the Use of Internet Exchange Points: A Guide to Policy, Management, and Technical Issues* (Tech. Rep.). Internet Society. Retrieved from <http://www.internetsociety.org/promoting-use-internet-exchange-points-guide-policy-management-and-technical-issues>
- Kende, M., & Hurpy, C. (2012, April). *Assessment of the Impact of Internet Exchange Points—Empirical Study of Kenya and Nigeria* (Report for the Internet Society No. 20945-144). Internet Society. Retrieved from <http://internetsociety.org/ixpimpact>