

Quantifying Deployability & Evolvability of Future Internet Architectures via Economic Models

Li Ye, Hong Xie, John C.S. Lui

Department of Computer Science and Engineering, The Chinese University of Hong Kong

Abstract—Emerging new applications demand the current Internet to provide new functionalities. Although many future Internet architectures and protocols have been proposed to fulfill such needs, ISPs have been reluctant to deploy these architectures. We believe technical issues are not the main reasons as many of these new proposals are technically sound. In this paper, we take an economic perspective and seek to answer: *Why most new Internet architectures failed to be deployed? What makes a new architecture easier to deploy?* We develop a game-theoretic model to characterize the outcome of an architecture’s deployment through the equilibrium of ISP’s decisions. We also use our model to explain the deploying outcomes of IPv6, DiffServ, CDN, etc., and the “Internet flattening phenomenon”. Furthermore, one can use our model to predict the deployability of new architectures such as NDN, XIA, and compare the deployability of competing architectures. Our study also suggests that the architectures which try to make a fresh start may have a low deployability, unless they have efficient incremental deployment mechanisms or one introduces a centralized coordinator to help the deployment.

I. INTRODUCTION

There is always a constant push for the Internet to be “*evolvable*” so as to support new applications with new functionalities. For example, huge amount of streaming video traffics from Netflix require highly efficient content delivery across the Internet. Also, users of online social network services like Facebook want their private chats to be securely protected. Moreover, the increasing number of mobile phones and IoT devices require better mobility and security support. However, many of these needs are not being supported by the IPv4 network. To support these emerging needs, researchers have been developing new architectures and protocols, and more importantly, exploring how to make the Internet “*evolvable*” so to incorporate new functionalities. Unfortunately, many of these research efforts fail to lead to wide scale deployment.

In 1990s, the protocol IPv6 was designed to improve IPv4. In particular, IPv6 aims to provide more addresses and new features such as security. However, after 20 years of effort, less than 20% of the current Internet traffic is using IPv6[1]. Differentiated service (DiffServ)[2] was designed to provide QoS guarantee. Although it is supported by many commercial routers[2], only few Internet service providers (ISPs) are willing to turn on the DiffServ functions. In contrast, the content delivery network (CDN)[3] technology enjoys a rapid growth of deployment. Now, over 50% of the Internet traffic is delivered by CDNs[3]. In the past decade, a number of future Internet architectures, e.g., NDN and XIA, were proposed and they wanted to make fundamental changes to the current

IP network. NDN[4] natively facilitates content distribution, while XIA[5] provides intrinsic security and enables incremental deployment of future technologies. MobilityFirst[6] aims to support mobile devices. Although NDN, XIA and MobilityFirst all have functional prototype systems, up to the writing of this paper, they still lack the wide-scale deployment.

All of the above architecture/protocol proposals claimed to improve the current Internet if they are successfully deployed. However, only the CDN technology is smoothly deployed in the Internet, while most of the others are not. This motivates us to explore: *Why most new Internet architectures failed to be deployed? What makes a new architecture/protocol easier to deploy?* It is interesting to note that the deployment failure of many proposed architectures/protocols is not due to the technical issue. In fact, many of these proposed architectures have superior designs than the present IPv4 network. Instead, we argue that “*economic issues*” are often crucial in deciding the deployability of any new Internet architecture or protocol.




Fig. 1: Economic difficulty to deploy a new Internet architecture

To illustrate, let us consider the following example. Fig. 1 depicts a simple network with three ISPs. There is a traffic flow from ISP 1 to ISP 3. Let’s say under the current Internet architecture, the whole network can gain a total revenue of \$20. Suppose a new Internet architecture, if it is fully deployed, will increase the total revenue of the whole network to \$32 (i.e., improve the revenue by \$12). Each ISP has a launching cost of \$3 to deploy this new architecture. Suppose the revenue improvement is evenly distributed among ISPs, i.e., each ISP gains $12/3 = \$4$, or each ISP will earn \$4 more by investing \$3 to upgrade the architecture, which yields a net gain of $\$4 - \$3 = \$1$. However, the deployment requires a “full-path” participation, i.e. ISP 1, ISP 2 and ISP3 all need to deploy, otherwise the functionality of this new architecture will not be enabled. Fig. 1 shows that when only ISP 1 and ISP 2 deploy (ISP 3 does not deploy), the total revenue will not be improved as the new functionality is not enabled.

Although the above example illustrates the potential profit gain for each ISP, unfortunately, the new architecture will not

necessarily be deployed in the network. The reason is that ISPs are not certain that they will reap any benefit if they alone decide to deploy the new architecture. In fact, if an ISP deploys, she will gain \$1 only if the other two ISPs will also deploy such architecture, otherwise the ISP will lose \$3. The main difficulty is that an ISP cannot be certain that other ISPs will make a similar decision on deployment. Given the uncertainty of others' participation, ISPs tend to be conservative and this will lead to the failure of deployment. The above example highlights that a new Internet architecture can be difficult to deploy even if it can bring higher profits.

This paper studies the economic issues for the deployment of new Internet architecture/protocol, and we aim to answer:

- *Why many new Internet architectures/protocols, e.g., IPv6, NDN, XIA, failed to receive large scale acceptance and wide deployment by ISPs?*
- *Under what conditions a new Internet architecture can be successfully deployed? How to compare multiple competing architectures and protocols?*

In addition, we analyze the economic impact of some engineering mechanisms, e.g., tunneling, that were proposed for the incremental deployment of new Internet architectures. We also study the “Internet flattening phenomenon”, where content providers are bypassing ISPs but instead, place their servers in data centers close to the end users. Our contributions are:

- We build economic models to quantify the condition for a new Internet architecture to be deployed by ISPs. Our model also allows us to compare the deployability of competing architectures and see which one will win.
- Our model indicates that a profitable new architecture may not be widely deployed, and a superior architecture may lose to another competing architecture. It explains why architectures like IPv6, DiffServ are difficult to deploy, while deploying CDN, NAT is easy.
- We quantify how incremental deployment mechanisms such as IPv6 tunneling improve the deployability of an architecture. Furthermore, we show by relying on data centers, content providers will be easier to deploy new architectures in the flattened Internet. We also show how an architecture that is originally not deployable could be deployed, by a centralized economic mechanism.

II. SYSTEM MODEL

A. Model of The Internet Network

We use an undirected graph $\mathcal{G} \triangleq (\mathcal{N}, \mathcal{E})$ to characterize the inter-connections between ISPs. The node set $\mathcal{N} \triangleq \{1, \dots, N\}$ denotes the set of all ISPs. Each node corresponds to one ISP, and physically consists of all the routers, switches, etc., operated by that ISP. Fig. 2 illustrates a set $\mathcal{N} = \{1, 2, 3, 4\}$ of four ISPs, where ISP 1 has one router, ISP 2 has three routers, etc. The edge set $\mathcal{E} \subseteq \{(i, j) | i, j \in \mathcal{N}, i < j\}$ indicates the connectivity among ISPs, where we require $i < j$ to eliminate the redundancy of the same undirected edges (i, j) and (j, i) . If two ISPs are inter-connected, then there is a physical link between them. For example, in Fig. 2, we have $\mathcal{E} = \{(1, 2), (2, 3), (2, 4)\}$.




Fig. 2: Flows via the routers




Fig. 3: Flows via the ISPs

We define a flow as an aggregation of all the traffic of packets along a route. Formally, we use the vector of all traversed ISPs along the route $\mathbf{f} \triangleq (s, \dots, t)$ to represent a flow, where nodes are ordered according to the direction of the flow such that $s \in \mathcal{N}$ denotes the source node and $t \in \mathcal{N}$ denotes the destination node. For example, Fig. 2 depicts the routing path of flow 1 over all the routers (or switches). Fig. 3 hides the details of routers, and shows that flow 1 goes through ISPs 3, 2, 4, 2 and 1, accordingly $\mathbf{f}_1 = (3, 2, 4, 2, 1)$. In this example, the flow has a twisted routes which goes through ISP 2 more than once. We allow such duplication of nodes in a flow \mathbf{f} to capture the full interactions among ISPs induced by network traffics. Let \mathcal{F} denote the set of all the Internet traffic flows. In Fig. 3, we have $\mathcal{F} = \{\mathbf{f}_1, \mathbf{f}_2\} = \{(3, 2, 4, 2, 1), (4, 2, 3)\}$. As one will see later, this flow representation facilitates the analysis of ISPs' economic behaviors. Note that our flow representation reflects its routing path rather than its traffic volume, while our model applies to flows with arbitrary traffic volume. Note that in reality, if some ISPs of a flow do not deploy the new protocol, then in the routing path, these ISPs might be replaced by other ISPs who deploy the new protocol. In this paper, we do not consider the change of routing path, because the change of routing path is not a reason why deploying a new architecture is difficult as shown in our technical report[7].

B. Model of Future Internet Architecture

Launching cost. Let $c_i \in \mathbb{R}_+$ denote the launching cost to deploy a new architecture, i.e. the money needed to start the new architecture. It captures the expenses to purchase new hardware, to upgrade software, and to pay for engineers who use the new architecture, etc. Some expenses to operate the new architecture occur immediately after the deployment, so we also include them in the launching cost. We assume that the upgrading is indivisible, i.e., an ISP either upgrades all its network or upgrades none. If an ISP divides its network into multiple sub-networks and each sub-network independently decides whether to deploy, one can regard each sub-network as an ISP in our model. In our technical report[7], we show that allowing such divisions of ISPs will not affect our results on whether an architecture will be successfully deployed.

Revenue Improvement. ISPs earn revenue from customers (e.g. end users or content providers). With new functionalities of the new architecture, the customers may be willing to pay more. We aim to quantify the improvement of revenue from the new architecture after the deployment. Let $R(\mathcal{H}, \mathcal{S})$ denote the total revenue generated from a set of flows $\mathcal{H} \subseteq \mathcal{F}$ when a set \mathcal{S} of ISPs deploy the new architecture. In particular, $R(\mathcal{H}, \emptyset)$ is the total revenue from flows \mathcal{H} when no ISP deploy the new

architecture, or, the total revenue from the current architecture.

Assumption 1. For all $\mathcal{H}_1, \mathcal{H}_2 \subseteq \mathcal{F}$, $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$, it holds that $R(\mathcal{H}_1 \cup \mathcal{H}_2, \mathcal{S}) = R(\mathcal{H}_1, \mathcal{S}) + R(\mathcal{H}_2, \mathcal{S})$.

Assumption 1 implies that the total revenue of the whole network \mathcal{G} is the summation of the revenue generated by each individual flow, i.e., $R(\mathcal{F}, \mathcal{S}) = \sum_{\mathbf{f} \in \mathcal{F}} R(\{\mathbf{f}\}, \mathcal{S})$.

The revenue of a flow depends on to what extend the functionality is enabled. The enabling of a new functionality depends on how many ISPs deploy the new architecture.

Definition 1. An ISP $i \in \mathcal{N}$ is a critical ISP to a flow \mathbf{f} , if its action (deploy or not) influences the functionality of the new architecture for the flow \mathbf{f} ; otherwise, it is non-critical.

Let $\mathcal{C}(\mathbf{f})$ denote a set of all critical ISPs to flow \mathbf{f} . Namely, to utilize the functionality of a new architecture on a flow \mathbf{f} , all ISPs in $\mathcal{C}(\mathbf{f})$ must deploy the new architecture. For example, consider flow \mathbf{f}_1 in Fig. 3. To utilize the functionality of a network-layer protocol like IPv6 in flow \mathbf{f}_1 , the set of critical ISPs is $\mathcal{C}(\mathbf{f}_1) = \{1, 2, 3, 4\}$. To utilize the functionality of a transport layer protocol like TCP, $\mathcal{C}(\mathbf{f}_1) = \{1, 3\}$. Finally, to have the caching functionality of a CDN service, we have $\mathcal{C}(\mathbf{f}_1) = \{1\}$. The reason is that TCP requires only the sender and the receiver to participate, and for the CDN, only the “last-mile” ISP that provides the CDN service needs to deploy.

Definition 2. A new architecture requires a full-path participation, if $\mathcal{C}(\mathbf{f})$ contains all ISPs in the flow \mathbf{f} , for all $\mathbf{f} \in \mathcal{F}$.

IPv6 requires full-path participation while TCP and CDN do not. Let $n(\mathbf{f}, \mathcal{S}) \triangleq |\mathcal{C}(\mathbf{f}) \cap \mathcal{S}|$ denote the number of critical ISPs of \mathbf{f} who participate in the deployment.

Incremental deployment mechanisms were developed to enable the functionality of a new architecture when not all critical ISPs participate in the deployment. A good example is “IPv6 tunneling”, where tunnels are built to bypass those IPv4 ISPs so that IPv6 packets could be delivered.

Assumption 2. Given $\mathcal{S}, \mathcal{T} \subseteq \mathcal{N}$, the function $R(\cdot, \cdot)$ satisfies:

- 1) $R(\{\mathbf{f}\}, \mathcal{S}) \leq R(\{\mathbf{f}\}, \mathcal{T})$ if $\mathcal{S} \subseteq \mathcal{T}$;
- 2) $R(\{\mathbf{f}\}, \mathcal{S}) = R(\{\mathbf{f}\}, \emptyset)$, if the new architecture has no incremental deployment mechanisms and $n(\mathbf{f}, \mathcal{S}) < |\mathcal{C}(\mathbf{f})|$.
- 3) $R(\{\mathbf{f}\}, \mathcal{S}) \geq R(\{\mathbf{f}\}, \emptyset)$, if we could apply incremental deployment mechanisms to the new architecture.

Assumption 2 captures that: (1) The revenue generated from a flow is non-decreasing as more ISPs deploy the new architecture; (2) Suppose the new architecture cannot be incrementally deployed, then the revenue from a flow will not be changed if any critical ISPs for this flow do not deploy the new architecture. (3) If there are incremental deployment mechanisms, then the new architecture partially works without complete deployment, and ISPs’ revenue might be improved.

Note that 2) in Assumption 2 may depart from reality. This assumption ignores the decline of revenue from the old architecture due to the competition of the new architecture. In reality, ISPs’ adoption of new architecture might hurt their revenue from the flows that use the old architecture. This reflects to the notion of “self-cannibalism”, which was thought to be a reason why ISPs are reluctant to deploy. In this paper, we

focus on the baseline case where $R(\{\mathbf{f}\}, \mathcal{S}) = R(\{\mathbf{f}\}, \emptyset)$. Our model can be extended[7] to show that the “self-cannibalism” is not a reason why a new architecture is difficult to deploy.

The performance with incremental deployment mechanisms depends on how many critical ISPs deploy the new architecture. Let us take the example in Fig. 3, where two ISPs 3 and 4 deploy IPv6 in the flow \mathbf{f}_2 (while ISP 2 does not). To make IPv6 feasible in such incremental deployment situation, IPv6 packets need to bypass ISP 2 which uses IPv4. In this flow, no matter which two ISPs deploy IPv6 (e.g. ISPs 2 and 3), a IPv6 packet needs to bypass one ISP via tunneling.

Assumption 3. Given $\mathcal{S}, \mathcal{T} \subseteq \mathcal{N}$, if $n(\mathbf{f}, \mathcal{S}) = n(\mathbf{f}, \mathcal{T})$, then $R(\{\mathbf{f}\}, \mathcal{S}) = R(\{\mathbf{f}\}, \mathcal{T})$.

Assumption 3 captures that the revenue generated from a flow is determined by the number of critical ISPs who deploy the new architecture. In fact, Assumption 3 will be automatically satisfied without incremental deployment mechanisms.

Assumption 4. For all $\mathcal{S} \subseteq \mathcal{T}$, $i \in \mathcal{N}$ and $i \notin \mathcal{T}$, it holds that $R(\mathcal{F}, \mathcal{S} \cup \{i\}) - R(\mathcal{F}, \mathcal{S}) \leq R(\mathcal{F}, \mathcal{T} \cup \{i\}) - R(\mathcal{F}, \mathcal{T})$.

Assumption 4 states that as more ISPs have deployed the new architecture, the marginal revenue gain by adding one more ISP to deploy the new architecture is larger. Namely, different ISPs are complementary to generate revenue in the deployment of the new architecture. Assumption 4 will be automatically satisfied without incremental deployment mechanisms. We denote ISPs’ revenue gain from deployment as

$$v(\mathcal{S}) \triangleq R(\mathcal{F}, \mathcal{S}) - R(\mathcal{F}, \emptyset), \quad \forall \mathcal{S} \subseteq \mathcal{N}.$$

Note that $v(\emptyset) = 0$. Moreover, we are only interested in the cases where the revenue gain $v(\mathcal{S})$ from deploying the new architecture is non-negative, as implied by our Assumption 2. Hence we also call the revenue gain the “benefit”. We next proceed to analyze ISPs’ non-trivial decisions when $v(\mathcal{S}) \geq 0$.

III. ISPs’ STRATEGIC BEHAVIOR

We characterize the “stable” mechanism for the ISPs to distribute the revenue gain. Then, we formulate a game to capture ISPs’ strategic behaviors in deploying the new architecture.

A. Revenue Gain Distribution Mechanism

Currently, ISPs distribute the revenues from the old functionality (i.e. the basic packets transmission) according to bilateral contracts and bilateral peering agreements. The existing contracts and peering agreements are not concerned in this paper. The unsettled distribution of our interest is on the revenue gain introduced by new functionalities of the new architecture. Let $\phi(\mathcal{S}, v) = (\phi_1, \dots, \phi_N) \in \mathbb{R}_{\geq 0}^N$ denote a distribution mechanism, where $\phi_i(\mathcal{S}, v)$ is the revenue gain distributed to ISP $i \in \mathcal{N}$.

Distributing the revenue gain is non-trivial, because ISPs may refuse to deploy the architecture if the revenue gain is not “fairly” distributed. A “stable” revenue gain distribution mechanism should satisfy the following four properties.

Property 1 (Efficiency). $\sum_{i \in \mathcal{S}} \phi_i(\mathcal{S}, v) = v(\mathcal{S})$.

It captures that all the revenue gain is only distributed to the participating ISPs, and nothing remains undistributed.




Fig. 4: Example Flows




Fig. 5: Multiple equilibria

Property 2 (Symmetry). If $i, j \in \mathcal{S}$ satisfies $v(\mathcal{T} \cup \{i\}) = v(\mathcal{T} \cup \{j\})$ for all $\mathcal{T} \subseteq \mathcal{S} \setminus \{i, j\}$, then $\phi_i(\mathcal{S}, v) = \phi_j(\mathcal{S}, v)$.

This property captures the “fairness”, i.e., if two ISPs have indistinguishable contributions, we have no reasons to distribute more money to one of the ISPs. In fact, these two ISPs have the same *bargaining power* during the distribution.

Property 3 (Dummy). If $i \in \mathcal{S}$ is a dummy ISP, i.e., $v(\mathcal{T}) = v(\mathcal{T} \cup \{i\})$ holds for all $\mathcal{T} \subseteq \mathcal{S} \setminus \{i\}$, then $\phi_i(\mathcal{S}, v) = 0$.

It captures that if an ISP has no contribution to the revenue gain, then this ISP is dummy and should receive nothing.

Property 4 (Additivity). For two revenue gain functions v and w , $\phi_i(\mathcal{S}, v+w) = \phi_i(\mathcal{S}, v) + \phi_i(\mathcal{S}, w)$ for any $i \in \mathcal{S}$.

It captures that the distribution of revenue gain from some new architecture should be consistent no matter whether it is composed with other orthogonal architecture or not.

As stated in economic theory[8], distributing the revenue gain according to the Shapley value is the “unique choice” to satisfy the above four properties. Formally, according to the Shapley value, each ISP $i \in \mathcal{S}$ receives a revenue gain of

$$\phi_i(\mathcal{S}, v) = \sum_{\mathcal{T} \subseteq \mathcal{S} \setminus \{i\}} \frac{|\mathcal{T}|!(|\mathcal{S}|-|\mathcal{T}|-1)!}{|\mathcal{S}|!} [v(\mathcal{T} \cup \{i\}) - v(\mathcal{T})]. \quad (1)$$

Besides the above axiomatic approach, it was also shown[9] that the Shapley values will be distributed if ISPs are free to do non-cooperative bargaining. Shapley value has been applied to study Network Economics problems[10]. Note that computing the Shapley value is NP-hard in general[11]. Under mild assumptions in Sec. II-B, we have a closed-form expression of Shapley value.

Theorem 1. Under Assumption 1 and 3, (1) reduces to

$$\phi_i(\mathcal{S}, v) = \sum_{\mathbf{f} \in \mathcal{F}} \mathbf{1}_{\{i \in \mathcal{C}(\mathbf{f})\}} \frac{R(\{\mathbf{f}\}, \mathcal{S}) - R(\{\mathbf{f}\}, \emptyset)}{n(\mathbf{f}, \mathcal{S})}, \forall i \in \mathcal{S}. \quad (2)$$

Remark. Due to the page limit, all proofs are in our technical report[7]. Eq. (2) states that the revenue gain of a flow \mathbf{f} is evenly shared among all $n(\mathbf{f}, \mathcal{S})$ critical ISPs that deploy the new architecture, and the Shapley value of an ISP is the total shared revenue gains from different flows.

To illustrate, let us consider Fig. 4, which depicts a network consisting of 4 ISPs and 6 flows where each flow passes through three ISPs. Suppose the new architecture requires a full-path participation and all four ISPs $\mathcal{S} = \{1, 2, 3, 4\}$ upgrade to the new architecture. Each flow has a revenue gain of \$3, i.e. $R(\mathbf{f}, \mathcal{S}) - R(\mathbf{f}, \emptyset) = \$3, \forall \mathbf{f} \in \mathcal{F}$. According to (2), each ISP shares $\$3/3 = \1 in each flow. In total, ISP 1, 3 and 4 gain \$4 because they participate in 4 flows. One can see that ISP 2

shares a higher Shapley value of \$6, because all 6 flows must go through it. In other words, ISP 2 has a higher contribution to the revenue gain of the new architecture. We like to point out that the Shapley value is calculated ex-post and not known to ISPs before deployment. Here, the Shapley values serve as the benchmarking “true” benefits that the ISPs need to estimate, while Shapley values may depart from realistic distributions.

B. The Architecture Deployment Game

Given the mechanism ϕ , the revenue gain share of each ISP is determined by her action (i.e. to deploy the architecture or not) and the other ISPs’ actions. We formulate a strategic-form game to characterize ISPs’ strategic behavior in deciding whether to deploy the new architecture. We denote $\tilde{\mathcal{C}} \triangleq \bigcup_{\mathbf{f} \in \mathcal{F}} \mathcal{C}(\mathbf{f})$ as the set of all critical ISPs. If an ISP is not critical to any flow, she is dummy and will not get involved in the deployment. Hence the players of interests are all the critical ISPs $\tilde{\mathcal{C}}$. Each ISP has two possible actions denoted by $\mathcal{A} \triangleq \{0, 1\}$, where 1 indicates that an ISP deploys the new architecture and 0 indicates not. Let $a_i \in \mathcal{A}$ denote the action of ISP $i \in \tilde{\mathcal{C}}$ and let $\mathbf{a} \triangleq (a_i)_{i \in \tilde{\mathcal{C}}}$ denote the action profile of all critical ISPs. Given an action profile \mathbf{a} , the corresponding set of ISPs who deploy the new architecture is denoted as $\mathcal{S}_{\mathbf{a}} \triangleq \{i | a_i = 1, i \in \tilde{\mathcal{C}}\}$. The utility (or profit) $u_i(\mathbf{a})$ of an ISP is the shared revenue gain minus its launching cost, i.e.

$$u_i(\mathbf{a}) \triangleq \begin{cases} \phi_i(\mathcal{S}_{\mathbf{a}}, v) - c_i & \text{if } a_i = 1, \\ 0 & \text{if } a_i = 0. \end{cases} \quad (3)$$

We denote this “architecture deployment game” by a tuple $G \triangleq (\tilde{\mathcal{C}}, \mathcal{A}, \mathbf{u})$, where $\mathbf{u} \triangleq (u_1, \dots, u_N)$ is a vector of functions.

Game G is a “potential game”, where ISPs have a “potential function” as the common objective during the deployment.

Lemma 1 ([12]). If ϕ satisfies (1), G is a potential game, i.e.,

$$u_i(1, \mathbf{a}_{-i}) - u_i(0, \mathbf{a}_{-i}) = \Phi(1, \mathbf{a}_{-i}) - \Phi(0, \mathbf{a}_{-i}), \quad (4)$$

holds for all \mathbf{a}_{-i} , some $\Phi: \mathcal{A}^{|\tilde{\mathcal{C}}|} \rightarrow \mathbb{R}$, where $\mathbf{a}_{-i} \triangleq (a_j)_{j \neq i}$.

Theorem 2. Under Assumption 1, 3, a potential function is:

$$\Phi(\mathbf{a}) = \sum_{\mathbf{f} \in \mathcal{F}} \sum_{m=1}^{n(\mathbf{f}, \mathcal{S}_{\mathbf{a}})} \frac{V(\mathbf{f}, m)}{m} - \sum_{i \in \tilde{\mathcal{C}}} a_i c_i, \quad (5)$$

where $V(\mathbf{f}, m) \triangleq R(\{\mathbf{f}\}, \mathcal{S}) - R(\{\mathbf{f}\}, \emptyset)$ with $n(\mathbf{f}, \mathcal{S}) = m$. We call $B(\mathcal{S}_{\mathbf{a}}) \triangleq \sum_{\mathbf{f} \in \mathcal{F}} \sum_{m=1}^{n(\mathbf{f}, \mathcal{S}_{\mathbf{a}})} \frac{V(\mathbf{f}, m)}{m}$ total immediate benefits.

Remark. Eq. (4) states that any ISP’s profit to deploy, is the same as the change of a unified potential function $\Phi(\mathbf{a})$. Thus, an ISP will have a positive profit to deploy an architecture if and only if her deployment increases the potential function.

This potential function $\Phi(\mathbf{a})$ has insightful physical meanings. The term $V(\mathbf{f}, m)/m$ is the revenue gain distributed to the m th deployer in flow \mathbf{f} immediately after her deployment. The term $\sum_{m=1}^{n(\mathbf{f}, \mathcal{S}_{\mathbf{a}})} V(\mathbf{f}, m)/m$ is all such immediate benefits that have been distributed to the past deployers in flow \mathbf{f} . Summing up on all flows, $B(\mathcal{S}_{\mathbf{a}})$ is the total benefits that have been distributed to the ISPs $\mathcal{S}_{\mathbf{a}}$ immediately when they deploy the new architecture. Also, $\sum_{i \in \tilde{\mathcal{C}}} a_i c_i$ is the total launching cost of ISPs $\mathcal{S}_{\mathbf{a}}$. Therefore, the potential function $\Phi(\mathbf{a})$ is the total immediate benefits minus the total launching cost.

IV. ANALYZING ISPs' DECISIONS VIA EQUILIBRIUM

An equilibrium represents a stabilized deployment status of the critical ISPs of a new architecture. We first show that multiple equilibria are possible in the architecture deployment game. Then, we show the equilibrium that maximizes the potential function will be reached when the ISPs face uncertainty.

A. Deployment Equilibria

Definition 3. An action profile $\mathbf{a}^* \in \mathcal{A}^{|\tilde{\mathcal{C}}|}$ is a pure strategy Nash equilibrium of the game G , if for any $i \in \tilde{\mathcal{C}}$, it holds that

$$u_i(\mathbf{a}_i^*, \mathbf{a}_{-i}^*) \geq u_i(\mathbf{a}_i, \mathbf{a}_{-i}^*), \quad \forall \mathbf{a}_i \in \mathcal{A}.$$

Namely, under such equilibrium, each ISP can not increase her utility by unilaterally deviating from her current action. There exists a one-to-one mapping between equilibria of the potential game and local maxima of the potential function[12].

Our game G may have multiple equilibria, as the potential function may have multiple local maxima. To illustrate, consider Fig. 5. There are two flows and upgrading to the new architecture will bring a benefit \$12 for each flow. The launching cost of each ISP is \$3. There are 4 equilibria. The first is illustrated in Fig. 5: (a) all ISPs do not deploy the new architecture. This is because an ISP's unilateral deviation to deploy the new architecture will result in a loss of \$3. Other three equilibria are: (b) ISPs 1,2,3 deploy and ISPs 4,5 do not deploy; (c) ISPs 2,4,5 deploy and ISPs 1,3 do not deploy; (d) All ISPs deploy. In this example, we notice that "No-deployment" (i.e., (a) all critical ISPs do not to deploy) can be an equilibrium, while the successful deployment (i.e., (d) all critical ISPs deploy) can also be an equilibrium. Among those equilibria, there is a smallest (or largest) equilibrium with the smallest (or largest) set of deployed ISPs in our game[13].

B. Which Equilibrium Will Be Reached?

We present two approaches to show eventually the equilibrium that maximizes the potential function will be reached.

(A) Logit response dynamics. We divide the time into slots, i.e., $t \in \{1, \dots, T\}$. Let $\mathbf{a}^{(t)}$ denote the action profile at time slot t , and ISPs start with some initial action profile $\mathbf{a}^{(0)}$. At time slot t , we randomly pick one ISP, let's say $i \in \tilde{\mathcal{C}}$, to make a decision based on other ISPs' actions in the last time slot, i.e. $\mathbf{a}_{-i}^{(t-1)}$. This setting captures that ISPs sequentially make decisions. More specifically, ISP i chooses each action $a_i^{(t)} \in \{0, 1\}$ with a probability that is logit-weighted by utility:

$$Pr[a_i^{(t)} = a | \mathbf{a}_{-i}^{(t-1)}] = \frac{e^{\beta_t u_i(a, \mathbf{a}_{-i}^{(t-1)})}}{e^{\beta_t u_i(0, \mathbf{a}_{-i}^{(t-1)})} + e^{\beta_t u_i(1, \mathbf{a}_{-i}^{(t-1)})}}, \quad a \in \{0, 1\}.$$

This logit-weighted choice gives ISP i the optimal expected utility, when the ISP has some uncertainty about his utility[14]. The parameter $\beta_t \geq 0$ represents the degree of uncertainty. Low β_t means high uncertainty. When $\beta_t = 0$, the ISP is the most uncertain and chooses to deploy with probability 0.5. When $\beta_t \rightarrow +\infty$, ISPs always choose the action with the highest utility and the *logit response* reduces to the *best response*. Since we allow ISPs to randomly choose to deploy or not, we look into the probability to observe each deployment status.

Lemma 2 ([15]). Suppose $\beta_t = \beta, \forall t$. For the game G , the limiting distribution of the logit response dynamics is

$$\mathbb{P}(\mathbf{a}) = e^{\beta \Phi(\mathbf{a})} / \sum_{\mathbf{a} \in \mathcal{A}^{|\tilde{\mathcal{C}}|}} e^{\beta \Phi(\mathbf{a})}. \quad (6)$$

Remark. The ISPs are more likely to stay in the deployment status \mathbf{a} with a higher potential value $\Phi(\mathbf{a})$. When β is sufficiently large, the ISPs will eventually reach an equilibrium that maximizes the potential function, i.e., $\arg \max_{\mathbf{a} \in \mathcal{A}^{|\tilde{\mathcal{C}}|}} \Phi(\mathbf{a})$.

(B) Iterative elimination of dominated strategies. Another perspective is that ISPs make decisions under incomplete information of the revenue gain. Before deploying, ISP i perceives a benefit $(1 + \varepsilon_i)(R(\{\mathbf{f}\}, \mathcal{S}_a) - R(\{\mathbf{f}\}, \emptyset))$ for the flow \mathbf{f} , where ε_i is a random variable with distribution \mathcal{D}_i that could be both negative and positive values, and different ε_i 's are independent. A negative (or positive) perception error ε_i means ISP i is pessimistic (or optimistic) about the new architecture. We define $\lambda_i \triangleq (1 + \varepsilon_i) \in \mathbb{R}$ as the perception scaling factor. The perceived benefit $\lambda_i \times (R(\{\mathbf{f}\}, \mathcal{S}_a) - R(\{\mathbf{f}\}, \emptyset))$ is only known to ISP i , while the distribution \mathcal{D}_i is known to all the ISPs. We denote the strategy of an ISP i as a function from the perception scaling factor to her action, i.e. $s_i(\lambda_i) : \mathbb{R} \rightarrow \{0, 1\}$.

To investigate in ISPs' strategies, we use the solution concept called "iterative strict dominance". The basic idea is that ISPs will not choose those actions which are known to have worse profit in expectation. For example, an ISP will not deploy IPv6 if she will lose money by deploying it.

Lemma 3 ([16]). Under Assumption 4, as the error distributions \mathcal{D}_i 's concentrate around zero, a unique strategy profile survives from the iterative elimination of dominated strategy. Moreover, ISPs' actions $(s_i^*(1))_{i \in \tilde{\mathcal{C}}} \in \arg \max_{\mathbf{a} \in \mathcal{A}} \Phi(\mathbf{a})$.

Remark. This lemma states that under the iterative elimination of dominated strategies, the equilibrium that maximizes the potential function will be reached as the perception errors concentrate to zero. Some economic experiments were carried out[17] that coincide with this equilibrium prediction.

To illustrate, let us go back to our first example in Fig. 1. To start with, the ISPs know that when $\lambda_i \leq 0.75$, $s_i(\lambda_i) = 0, \forall i$. Namely, an ISP will not deploy if she perceives a total benefit which is less than \$9, or a benefit less than the launching cost \$3 for each ISP. For simplicity, suppose the distribution of the perception error \mathcal{D}_i is Gaussian with mean 0 and standard deviation 0.3 for $i=1,2,3$. Therefore, ISP 1 knows $\frac{\lambda_i}{\lambda_1} = \frac{(1 + \varepsilon_i)}{(1 + \varepsilon_1)}$ follows a ratio distribution of two independent Gaussian variables. Let us now show how ISP 1 determines her action for some $\lambda_1 > 0.75$. ISP 1 knows that some other ISP $i \in \{2,3\}$ will perceive a scaling factor $\lambda_i < 0.75$ with probability $P(\lambda_1) \triangleq Pr[\lambda_i = \lambda_1 \frac{(1 + \varepsilon_i)}{(1 + \varepsilon_1)} < 0.75]$. It indicates that some other ISP will not deploy with probability of at least $P(\lambda_1)$. In fact, when $\lambda_1 \leq 1.1$, the expected utility of ISP 1 is at most $\lambda_1 4(1 - P(\lambda_1))^2 - 3 \leq 0$, so ISP 1 will not deploy. In other words, $s_1(\lambda_1) = 0$ if $\lambda_1 \leq 1.1$. Now, we know that ISP 1 will not deploy when the perceived revenue gain is less than $\$12 \times 1.1 = \13.2 . Similar reasoning applies to other ISPs. Even without further inductions, the ISPs will choose not to deploy, when the perceived benefit is close to the true value \$12.

We have shown that the equilibrium which maximizes the potential function will be reached if an ISP is uncertain about benefit of new architecture and other ISPs' decisions, from two perspectives: (A) ISPs have dynamics by logit-weighted probabilities; (B) ISPs iteratively eliminate dominated strategies. Thus, we define this equilibrium as the “*robust equilibrium*”. In this paper, we focus on the baseline case of risk-neutral ISPs who maximize their expected utility. Our model can be extended to discuss realistic “risk-averse” or “risk-seeking” ISPs. For risk-averse ISPs who are unlikely to take the risk to deploy, the new architectures becomes even less deployable.

V. QUANTIFYING THE DEPLOYABILITY

Given a new architecture, we now can use our model to analyze ISPs' actions on whether to deploy it or not.

A. General Analysis of Deployability

Definition 4. An architecture is successfully deployed if in the “*robust equilibrium*”, all critical ISPs \tilde{C} choose to deploy.

We next define a “profitable” architecture, whose benefit can cover the total launching cost of all critical ISPs \tilde{C} .

Definition 5. An architecture is profitable if

$$v(\tilde{C}) \geq \sum_{i \in \tilde{C}} c_i. \quad (7)$$

It is a simple necessary condition for successful deployment. However, as we will see from a more refined necessary condition, some profitable architecture may not be deployed.

Corollary 1. An architecture is successfully deployed only if

$$B(\tilde{C}) \geq \sum_{i \in \tilde{C}} c_i \quad (8)$$

Condition (8) implies condition (7), but the reverse is not true.

Remark. This corollary comes from the requirement of $\Phi((1, \dots, 1)) \geq 0$ for “all critical ISPs to deploy” to be a robust equilibrium. It shows why a “profitable” architecture may not be successfully deployed. To illustrate, consider a network of three ISPs connected in a line topology. There is only one flow and all three ISPs are critical. An architecture has a total benefit which is twice of the total launching cost, i.e. $v(\tilde{C}) = 2 \sum_{i \in \tilde{C}} c_i$. Then $B(\tilde{C}) = \frac{1}{3} v(\tilde{C}) < \sum_{i \in \tilde{C}} c_i$, which violates condition (8) for successful deployment. Interestingly, even when the total benefit is twice of the total launching cost, the new architecture still cannot be successfully deployed.

Corollary 2. If condition (8) holds, then in the robust equilibrium, a non-empty set of ISPs will deploy the new architecture.

This corollary states that condition (8) is sufficient to guarantee that at least some of the ISPs (if not all) will deploy the new architecture. Now, we could see that condition (8) is a criteria to determine whether an architecture is deployable.

B. Impact of the Benefit-cost Ratio

Put $v(\tilde{C})$ as the common numerator, and (8) is equivalent to

$$\frac{v(\tilde{C})}{\sum_{i \in \tilde{C}} c_i} \geq \frac{v(\tilde{C})}{B(\tilde{C})} \triangleq \gamma. \quad (9)$$

Note that γ is defined as the ratio between the total benefits $v(\tilde{C})$ and the “total immediate benefits” $B(\tilde{C})$. From Corollary 1, we have $\gamma \geq 1$. Condition (9) says that the ratio between the total benefit and total launching cost $v(\tilde{C}) / \sum_{i \in \tilde{C}} c_i$ (“benefit-cost ratio” in short) should be higher than a “threshold” γ , for a new architecture to be deployable. For a fixed γ , if the benefit-cost ratio is higher, then the new architecture will have a higher chance to be deployed. Hence, the designer of an architecture should increase the benefit and reduce the launching cost. As we will show, γ reflects the properties of an architecture, and is related to the number of critical ISPs.

C. Impact of the Number of Critical ISPs

With no incremental deployment mechanisms, we have

$$\gamma = \frac{v(\tilde{C})}{B(\tilde{C})} = \frac{\sum_{\mathbf{f} \in \mathcal{F}} V(\mathbf{f}, |\mathcal{C}(\mathbf{f})|)}{\sum_{\mathbf{f} \in \mathcal{F}} \frac{V(\mathbf{f}, |\mathcal{C}(\mathbf{f})|)}{|\mathcal{C}(\mathbf{f})|}}. \quad (10)$$

This equation is a consequence of $v(\tilde{C}) = \sum_{\mathbf{f} \in \mathcal{F}} V(\mathbf{f}, |\mathcal{C}(\mathbf{f})|)$ and $B(\tilde{C}) = \sum_{\mathbf{f} \in \mathcal{F}} \frac{V(\mathbf{f}, |\mathcal{C}(\mathbf{f})|)}{|\mathcal{C}(\mathbf{f})|}$. In (10), γ is the harmonic mean of the number of critical ISPs $|\mathcal{C}(\mathbf{f})|$ weighted by the maximum benefit of each flow. The number of critical ISPs $|\mathcal{C}(\mathbf{f})|$ represents the “degree of coordination” required by the new architecture for flow \mathbf{f} . Then the physical meaning of γ is the “average degree of coordination” over the whole network.

On one hand, when an architecture requires a small number of critical ISPs for each flow, γ is small. Then, we observe from (9) that the architecture is easy to deploy if the benefit-cost ratio surpasses this small number. On the other hand, when a new architecture has a large number of critical ISPs for each flow, it will be difficult to deploy, since the required benefit-cost ratio is high. Let us see some real-world cases.

(1) Deployment difficulty of DiffServ: To have QoS guarantees offered by DiffServ, all ISPs along the path are critical. If any ISP along the path decides not to participate, the QoS is no longer guaranteed. Using our model, we know DiffServ requires a high benefit-cost ratio, and thus is difficult to deploy. Hence, we see little adoption of DiffServ even if QoS guarantee is urgently needed in the current Internet.

(2) The Internet flattening phenomenon: We are witnessing a flattening Internet[18], [19]. This happens as large content providers such as Google and Facebook use data centers near the end-users. Hence, the routing paths become shorter and many intermediate ISPs are bypassed. Interestingly, the flattening topology also brings about a more evolvable Internet. For many new architectures which requires full-path participation, the flattening Internet reduces the number of critical ISPs, which makes the architectures more deployable. In fact, with the prevalence of the data centers, many Internet flows are going through the data centers of a single ISP (or content provider). The intra-data center flows have only one critical node, which indicates the highest deployability according to our model. This explains why we see many innovations for data centers proposed and deployed. We have shown that by relying data centers, it will be easier for content providers to deploy new protocols/architectures that would not be deployed in the Internet. In this way, one can see that data centers help

content providers gain extra profit from the deployment of new protocols/architectures.

D. Impact of Incremental Deployment

With incremental deployment mechanisms, we have

$$\gamma = \frac{v(\tilde{\mathcal{C}})}{B(\tilde{\mathcal{C}})} = \frac{\sum_{\mathbf{f} \in \mathcal{F}} V(\mathbf{f}, |\mathcal{C}(\mathbf{f})|)}{\sum_{\mathbf{f} \in \mathcal{F}} \left(\sum_{m=1}^{|\mathcal{C}(\mathbf{f})|} \frac{V(\mathbf{f}, m)}{|\mathcal{C}(\mathbf{f})|} \right)}. \quad (11)$$

The incremental mechanisms enable the new architecture even when some of the critical ISPs do not deploy. The benefit from incremental deployment is reflected by $V(\mathbf{f}, m)$ in (11) where $m < |\mathcal{C}(\mathbf{f})|$. Compared to (10), it is clear that the incremental benefits brought by these mechanisms reduce the ratio γ , as the denominator in (11) becomes larger. This finding highlights the importance of these incremental deployment mechanisms for the final deployment of the new architectures.

(1) Incremental deployment mechanisms of IPv6: Different incremental deployment mechanisms[20] enable IPv6 in the current Internet by selecting ingress/egress points to bypass the non-IPv6 areas. Despite many of these mechanisms, almost all the IPv6 traffic are using native IPv6[1], which means these mechanisms are mostly not used. Based on this fact, we speculate that the ISPs do not have revenue gain by using these incremental mechanisms, i.e. $V(\mathbf{f}, m) = 0$ when $m < |\mathcal{C}(\mathbf{f})|$. In fact, if our speculation is wrong, many ISPs will use these mechanisms to improve their revenue. According to our model, the deployability of IPv6 will be as bad as if there were no incremental deployment mechanisms, since ISPs' incremental benefits is zero. In a word, these mechanisms are unsuccessful.

(2) XIA: XIA[21] is a future Internet architecture proposed recently that aims for an evolvable Internet. XIA has an intent-fallback system. If routers cannot operate on the primary “intent”, “fallbacks” will allow communicating parties to specify alternative actions. However, even if one has the best-possible incremental deployment mechanism, the incremental benefits are limited due to the characteristics of the aimed functionality. For example, it is almost impossible to have QoS guarantee without the participation of any ISP along the path, no matter which incremental deployment mechanism we use. Moreover, our model predicts that XIA itself is very hard to be deployed. This is because XIA is a network-layer protocol which has a large number of critical ISPs. Also, how much benefit it could bring to ISPs is not addressed.

VI. COMPETING ARCHITECTURES

Our results thus far consider one architecture only. In this section, we extend our model to study competing architectures with similar functionalities. We will show that a more “deployable” architecture will have a competitive advantage.

Deployment price of an architecture. ISPs charge customers for using the new functionality (e.g. CDN[22], DDos protection[23]). We consider a usage-based charging scheme. When all critical ISPs deploy the architecture, the unit price for the new functionality is $p \in \mathbb{R}_+$. Then, the revenue gain of a flow \mathbf{f} is $R(\{\mathbf{f}\}, \tilde{\mathcal{C}}) - R(\{\mathbf{f}\}, \emptyset) = p \times W(\mathbf{f})$, where $W(\mathbf{f})$ denotes the usage volume of the traffic flow \mathbf{f} . In compliant




Fig. 6: Competing architectures

with our previous model, the unit price for flow \mathbf{f} when a subset $\mathcal{S} \subset \tilde{\mathcal{C}}$ of ISPs deploy is $\frac{R(\{\mathbf{f}\}, \mathcal{S}) - R(\{\mathbf{f}\}, \emptyset)}{R(\{\mathbf{f}\}, \tilde{\mathcal{C}}) - R(\{\mathbf{f}\}, \emptyset)} p$. Then, the total immediate benefit is $B(\tilde{\mathcal{C}}) = (\sum_{\mathbf{f} \in \mathcal{F}} p \times W(\mathbf{f})) / \gamma$.

We define the deployment price of an architecture as the minimum unit price such that the condition (8) is satisfied, i.e., $p_d \triangleq \frac{\gamma \sum_{i \in \mathcal{C}} c_i}{\sum_{\mathbf{f} \in \mathcal{F}} W(\mathbf{f})}$. Namely, ISPs will deploy the architecture when the unit price is above p_d . Notice that the deployment price of an architecture depends on “degree of coordination” γ (about number of critical ISPs) and the total launching cost.

Multiple architectures under competition. Fig. 6 is an extension of the example in Fig. 1, where two architectures provide the same new functionality. Remember that architecture B requires all the ISPs 1,2,3 to deploy, and the launching cost for each ISP is \$3. Architecture A, on the other hand, requires only ISP 4 to deploy which is the ISP closest to the end user, and has a launching cost \$9 for that ISP. Note that the total launching costs for these two architectures are both \$9. For simplicity, we again assume no incremental deployment mechanism and we have one unit usage volume. *What will be a reasonable price for the new functionality?* As we can see, only if the price of the new functionality is higher than the “deployment price” \$27, the ISPs will deploy architecture B. Meanwhile, architecture A will be deployed when the price is higher than \$9. We argue that a customer will not pay as high as \$27 if he could enjoy the same functionality with a lower price \$9. Then, the customer will choose architecture A, and architecture B will not be deployed. We observe that architecture B fails to be deployed because the price is brought down by the more evolvable and competitive architecture A.

Suppose we have K new architectures providing the same functionality, with deployment prices $p_d^{(1)}, \dots, p_d^{(K)}$ respectively. We consider a market where ISPs are highly competitive so the customers can dictate the price. Here, we claim without rigorous proof that the customers will set the price to the lowest deployment price of these architectures, i.e. $\min_k \{p_d^{(k)}\}$. This is reasonable because the customers will not pay a higher price for a functionality if they could enjoy the same functionality with a lower price. Consequently, other architectures with higher deployment prices will not be successfully deployed. According to the deployment price, when competitive architectures have comparable total launching costs, the architecture with the lowest “degree of coordination” γ will win. Let us consider the following two cases.

(1) IPv6 vs. NAT: IPv6 and NAT (Network Address Translation) have a similar functionality of “addressing hosts”. The fact is that NAT is now supported by almost all Ethernet without enough IP addresses, while IPv6 is still not deployed in many countries. In short, NAT wins and this observation

could be explained by our model. IPv6 is a network-layer protocol which requires full-path participation. Although there are incremental deployment mechanisms, as we discussed before, they are rarely enabled by many ISPs. NAT is an application-layer solution which only requires end users who want more addresses to deploy. The average AS path length is around 4 [24]. Therefore, IPv6 requires higher “degree of coordination” (around four times) compared to NAT. In addition, we believe NAT has a much lower total launching cost, because IPv6 wants to change every router in the Internet. By definition, we know IPv6 has a higher (at least four times) deployment price than NAT. That is why we see NAT wins. In fact, ISPs seldom charge for IPv6 addresses, because the price of “addressing hosts” is brought down by NAT.

(2) NDN vs. CDN: Content Delivery Network (CDN) caches data spatially close to end-users to provide high availability and better performance. Meanwhile, Named Data Networking (NDN) [4] is a future Internet architecture that names data instead of their locations. Their main functionalities are both to provide scalable content delivery. CDNs operate at the application-layer, so only the CDN owners need to deploy the CDN infrastructures. NDN works directly at network layer[4]. Thus, the deployment of NDN requires full-path participation (or high degree of coordination and it is around four). Suppose NDN has a comparable total launching cost as CDN. Then we could see that NDN has a higher deployment price (around four times) than CDN, hence our models predict that the current design of NDN cannot be deployed and CDN will win out. Therefore, if NDN wants to be deployed faced with CDN’s competition, one needs to have efficient incremental deployment mechanisms. Guided by our analysis, a possible incremental deployment mechanism for NDN is to “run NDN over IP as edge caching”[25] at its early stage of deployment, which provides incremental benefits even when few ISPs deploy. This is possible since NDN is designed as a “universal overlay”. This simple modification might make NDN evolvable and competitive with CDN.

(3) Multipath TCP vs. Multipath QUIC: Multipath-TCP (MPTCP)[26] is an extension of TCP, which enables inverse multiplexing of resources, and thus increases TCP throughput. MPTCP requires the middleboxed (e.g. firewall) in the Internet to upgrade and not to interfere its packets[27]. Therefore, the critical nodes for MPTCP include the senders, receivers, and the ISPs with middleboxes. In contrast, Multipath-QUIC (MPQUIC)[28] is an extension of QUIC[29] that achieves similar functionalities of MPTCP. Because QUIC encrypts its packets and headers, MPQUIC avoids the interference from middleboxes. Then the critical nodes of MPQUIC only include the senders and receivers. We argue that the total launching cost of MPQUIC is not more than that of MPTCP. This is because MPTCP and MPQUIC both require the senders and receivers to upgrade their software, but MPTCP additionally requires the middleboxes to be upgraded. Also, MPQUIC has lower degree of coordination. Comparing the total launching cost and the degree of coordination, our models predict that MPQUIC will be deployed instead of MPTCP.

VII. ECONOMICS MECHANISM DESIGN

We observe that the difficulty of deployment comes from the “requirement of coordination” for decentralized ISPs. We then design a centralized mechanism to mitigate this difficulty. Our mechanism has the following two steps.

- 1) **Quoting:** Each ISP $i \in \tilde{C}$ submits a quote $q_i \in \mathbb{R}_+$ to the coordinator. An ISP’s quote implies a contract that the ISP would deploy the architecture once someone pays more than the quote. Quoting itself is not charged.
- 2) **ISP selection:** In this step, the coordinator selects a set of ISPs to deploy the architecture, and announces a reward for each of them. For a selected ISP, the announced reward should be at least as high as her quote. Then, the selected ISPs deploy the new architecture, and the coordinator gives ISPs the announced reward.

The coordinator in the above mechanism can be anyone such as the International organizations and the governments. In the second step of the mechanism, the coordinator selects the ISPs corresponding to the optimal solution $\mathcal{S}^*(\mathbf{q})$ the following problem (12), based on ISPs’ quotes $\mathbf{q} \triangleq (q_i)_{i \in \tilde{C}}$.

$$\begin{aligned} & \text{maximize}_{\mathcal{S} \subseteq \tilde{C}} && |\mathcal{S}|, \\ & \text{subject to} && \phi_i(\mathcal{S}, v) \geq q_i, \forall i \in \mathcal{S}. \end{aligned} \quad (12)$$

In addition, each selected ISP $i \in \mathcal{S}^*(\mathbf{q})$ gets a reward $\phi_i(\mathcal{S}^*(\mathbf{q}), v)$ which is at least as high as her quote.

Properties of the design. The launching cost is an ISP’s private information which the coordinator does not know. The ISPs may intentionally quote lower to increase the chance to be selected, or quote higher to ask for more reward. Our mechanism enforces ISPs to quote exactly the launching cost.

Theorem 3 (Truth-telling). *In our ISP selection mechanism, quoting $q_i = c_i$ is a weakly dominant strategy for each ISP $i \in \tilde{C}$.*

Theorem 4 (Efficiency). *The unique selection $\mathcal{S}^*(\mathbf{c})$ yields a maximal total revenue gain $v(\mathcal{S}^*(\mathbf{c}))$ of all ISPs and a maximal utility $u_i(\mathcal{S}^*(\mathbf{c}), v)$ for each individual ISP $i \in \tilde{C}$.*

Reducing the number of selected ISPs. Although the Internet can evolve provided a proper coordinator, a major difficulty is that the coordinator lacks the authority to manage a large number of ISPs. To make our mechanism practical, we select a small set of seeding ISPs called the “tipping set”[30], so that the system will go to another equilibrium with more deployed ISPs. After several rounds of tipping set selection, the ISPs will reach the largest equilibrium. Due to the page limit, we present the details of tipping set selection in our technical report[7] and show experiment results in Sec. VIII-C.

From the historical data[1], [31] for the transition from IPv4 to IPv6, we observe that actions of coordinators are highly correlated to IPv6’s deployment. Before the first World IPv6 Launch Day organized by Internet Society in 2012[32], less than 1% of users accessed their services over IPv6[1]. In 2018, this number goes to nearly 25%[33]. As another example, the Indian government decided a roadmap of IPv6’s deployment in July 2010[34] when the adoption rate is less than 0.5%. Now, over 30% of the traffics in India use IPv6[1]. In contrast, the government of China did not announce a plan

to put IPv6 into large-scale use until Nov. 2017[35]. Now, less than 3% of traffics in China use IPv6[1]. We argue that the different actions of Indian and Chinese governments determine the different deployment statuses of IPv6 in the two countries.

VIII. NUMERICAL EXPERIMENTS

A. Experiment Settings

Datasets. The first dataset[36] was collected from a European education network GÉANT with 23 ASes (Autonomous System). The data contains a network topology $G=(\mathcal{N},\mathcal{E})$ and a traffic matrix $\mathbf{T}\in\mathbb{R}_{>0}^{N\times N}$, where T_{ij} records the traffic volume from the source node i to the destination node j . There are 477 flows of non-zero traffic volume in this dataset. The second dataset[37] is the AS-level IPv4 topology collected by Caida on Dec. 2017. The dataset contains a weighted graph of 28,499 ASes $G=(\mathcal{N},\mathcal{E})$, where each edge $(i,j)\in\mathcal{E}$ can be either a direct or an indirect link from i to j . This dataset does not contain the traffic matrix data. Thus, we synthesize a traffic matrix based on the Gravity method[38]. The idea is that the traffic volume from node i to j is proportional to the *repulsive* factor of the source node i denoted by $T^{out}(i)$, and the *attractive* factor of the destination node j , denoted by $T^{in}(j)$, i.e. $T_{ij}\propto T^{out}(i)\times T^{in}(j)$. We apply the Clauset-Newman-Moore method[39] to extract the largest cluster in the network. This cluster contains 2,774 nodes, and we take $T^{in}(i)$ and $T^{out}(j)$ to be i.i.d. exponential random variables with mean 1. From 2774×2774 possible flows, we randomly select 74,424 (2%) as the flows with a positive demand of traffic for the new architecture.

Parameter settings. An ISP corresponds to an AS that has its own network policies, so we regard the ASes in the datasets as the ISPs in our model. It is known that GÉANT network uses IS-IS protocol[36] that implements the Dijkstra shortest path algorithm. For flows with a positive traffic, let $\mathcal{F}=\{\mathbf{f}$ is a shortest path from s to $t|s,t\in\mathcal{N},T_{s,t}>0\}$. As we discussed in Sec. VI, a new functionality such as CDN typically charges customers based on the usage. Again, we assume that the benefit of a flow is proportional to the usage volume and the unit price is p , i.e., $R(\{\mathbf{f}\},\tilde{\mathcal{C}})-R(\{\mathbf{f}\},\emptyset)=p\times W(\mathbf{f})$. Given a traffic matrix \mathbf{T} , the usage volume of a flow $\mathbf{f}=(s,\dots,t)$ is set as $W(\mathbf{f})=T_{s,t}$. The launching cost of an ISP depends on the workload of the ISP. Hence, we assume that the launching cost of an ISP i is proportional to the total amount of traffic through this ISP, i.e. $c_i=C\times\sum_{\mathbf{f}\in\mathcal{F}}\mathbf{1}_{\{i\in\mathbf{f}\}}W(\mathbf{f})$, where C is the launching cost for a unit amount of traffic. For an architecture deployment game $\langle\tilde{\mathcal{C}},\mathcal{A},\mathbf{u}\rangle$, as we scale p and C linearly at the same rate, the utility function \mathbf{u} scales linearly as well. This means that the Nash equilibria and the robust equilibrium will not be changed. Without loss of generality, we set $C=1$, and see the impact of p . Now, the parameter p also represents p/C .

B. Quantifying the Deployability

Benefit-cost ratio. The unit price p determines the benefit-cost ratio of the new architecture. For GÉANT network, the total benefit will be more than the total launching cost when $p>3.3$.

Fig. 7 shows the impact of p on the deployability without any incremental deployment mechanism. When $p\leq 4$, there is only one equilibrium that is “no ISP will deploy”, because the benefit of the new architecture is not enough to cover ISPs’ launching cost. When $p\geq 5$, the largest equilibrium (with the largest set of deployers) is that “all the 23 ISPs deploy”. However, in the robust equilibrium, no more than one ISP will deploy until $p\geq 10$. This is because the largest equilibrium will not have a positive potential function unless $p\geq 9.54$. For the IPv4 network with 2,774 ISPs, computing the robust equilibrium is intractable, so in Fig. 8, Fig. 10, and Fig. 12, we choose one of the smallest/largest equilibria with a higher potential function which is the one that is possible to be the robust equilibrium. For the IPv4 network, a new architecture will be profitable when $p\geq 4.78$. But only when $p\geq 21.88$, the condition (9) for successful deployment can be satisfied. Hence, as seen from Fig. 8, only when $p\geq 25$, “all ISPs to deploy” is the one that is possible to be a robust equilibrium. Similar phenomena are observed in both networks.

Incremental deployment mechanisms. We set the incremental benefit for \mathbf{f} as $R(\{\mathbf{f}\},\mathcal{S})-R(\{\mathbf{f}\},\emptyset)=\left(\frac{n(\mathbf{f},\mathcal{S})}{|\mathcal{C}(\mathbf{f})|}\right)^\alpha pW(\mathbf{f})$, when a set $\mathcal{S}\subseteq\tilde{\mathcal{C}}$ of ISPs deploy. The parameter $\alpha\geq 1$ represents performances of incremental deployment mechanisms, where a smaller α indicates better performances. When $\alpha=+\infty$, there is no incremental benefit. In this case, as depicted in Fig. 9, the new architecture will not be deployed in the GÉANT network until $p\geq 10$. In contrast, when $\alpha=1$, the architecture will be immediately deployed by 7 ISPs when $p=3$ and will be fully deployed by all ISPs when $p\geq 5$. As α decreases, the new architecture gets deployed by more ISPs for a fixed p , in help of better incremental mechanisms. As seen in Fig. 10, better incremental mechanisms also improve the deployability of new architectures in the IPv4 network.

Internet flattening phenomenon. To see the impact of a flattening Internet, we shrink the paths of the original flows to have a maximum length of M . For a original flow (s,v_1,\dots,v_L,t) with a path length $L+2>M$, the flattened flow will be $(s,v_{L-M+3},\dots,v_L,t)$ that contains s,t and $M-2$ ISPs which are nearest the destination t . This setting emulates that the sender uses data centers near the receiver. For GÉANT, Fig. 11 shows that when the maximum path length is shortened to $M=2$ (i.e. only the sender & receiver are in the flow), more than 10 ISPs will deploy when $p\geq 3.5$. Meanwhile, in the original network, ISPs will deploy the new architecture only when $p\geq 10$. Generally, the new architecture will be deployed by more ISPs in a more flattened network for a fixed p , as we can also see in Fig. 12 for the IPv4 network. One may observe that when content providers use data centers which are close to end users, they can pay a lower unit price p to the ISPs so to enjoy the deployed new architectures/technologies.

Lessons learned. A profitable new architecture may not be deployed. Also, an architecture with a higher benefit-cost ratio has higher deployability. The enhancement of incremental deployment mechanisms and the flattening Internet both improve the deployability of the new architectures.




Fig. 7: Unit price p and the scale of equilibria (GÉANT)




Fig. 8: Unit price p and the scale of equilibria (IPv4-net)




Fig. 9: Incremental deployment mechanisms (GÉANT)




Fig. 10: Incremental deployment mechanisms (IPv4-net)




Fig. 11: Impact of the flattening of network (GÉANT)




Fig. 12: Impact of the flattening of network (IPv4-net)




Fig. 13: Centralized economics mechanism (GÉANT)




Fig. 14: Logit-response dynamics (GÉANT)

C. Benefits of Economics Mechanisms

How the centralized coordinator can select the small “tipping set” of ISPs in the GÉANT network is shown in Fig. 13. At each time step $t=1, 2, \dots$, the coordinator selects some ISPs, where we see a vertical increment for the number of deployers. Between time steps t & $t+1$, some ISPs will deploy if they have higher utilities, where we see an increasing slope. For a small $p=3$, the economics mechanism cannot help and “all ISPs do not deploy” is the only equilibrium. When $p=7$, originally the new architecture cannot be successfully deployed as shown in Fig. 7. With the help of economics mechanism, by selecting 15 ISPs as the tipping set, all ISPs will finally deploy. When p further increases to 11, the coordinator only needs to select 7,2,4 ISPs in three steps respectively.

For the IPv4 network, when $p=5$, no ISP will deploy even with the existence of the coordinator. When $p=10$ and 20, after the coordinator selects 1,970 and 1,350 tipping ISPs respectively, all ISPs will finally deploy. Note that $1,350 < 1,970$. **Lessons learned.** As the benefit-cost ratio of architectures increases, the coordinator selects a smaller number of ISPs.

D. Logit-response Dynamics of ISPs

We simulate ISPs’ behaviors by the logit-response dynamics defined in Sec. IV-B, where we randomly initiate an ISP to deploy with probability 0.5. For the GÉANT network, we set $\beta_t = 8 \times 10^{-5} / t$, and take the average of 200 runs. As shown in Fig. 14, the number of deployed ISPs is close to the predictions of the robust equilibrium. When $T=50$, each ISP on the average makes two decisions, and the outcome of dynamics is very close to the robust equilibrium. As we increase $T=500$, the outcome becomes closer to the robust equilibrium. Similar results are observed for the IPv4 network[7]. The logit-response dynamics lead to the “robust” equilibrium, as if ISPs are maximizing some potential function in the deployment.

IX. RELATED WORKS

Designing future Internet architecture has been on the agenda since the early ages of the Internet[40]. A variety of future architectures were proposed[2], [4], [6], [5] to improve IPv4. Unfortunately, most of these proposals fail to deploy at scale. To make the Internet architectures evolvable, incremental deployment mechanisms are developed to enable universal access of IPv6[41], [42], [20]. While an evolvable new architecture should be compatible with old architectures, our work shows via economic models that an evolvable architecture should also provide incremental “benefits” to ISPs. Internet flattening phenomenon was studied in [18], [19], and our work formalizes their observations. A recent work [43] studied the incremental deployment of routing protocols, and suggested a coordinated adoption of a large number of ISPs.

Economics issues with the future Internet architectures have also been noticed. Wolf et al. developed ChoiceNet [44] to provide an economics plane to the Internet and a clear economics incentive for ISPs. Our work also points out that a new architecture may not be deployed even if it could be profitable for all ISPs. Ratnasamy et al.[45] has a similar “chicken-and-egg” argument. Our economics analysis strengthens these arguments and quantitatively analyze the difficulty of coordination among decentralized ISPs. Some works studied the adoptability of BGP security protocols[46][47]. They conducted simulations based on assumptions of ISPs’ behaviors, while we provide game-theoretic analysis to reveal key factors for the deployability, e.g. the coordination of ISPs. How to select some ISPs to stimulate the deployment was studied[48], but the incentives for these ISPs remain a problem. Our economic mechanism requires the coordinator to give nothing to the selected ISPs.

The “coordination failure” phenomenon was also studied in economics[49]. Monderer et al.[12] found that the equilibrium that maximizes a potential function accurately predicts Huyck’s experiments[50]. Then Morris et al.[51] give reasons

via the “global game”, which is used in our analysis.

X. CONCLUSION

This paper studies the deployability & evolvability of new architectures/protocols from an economic perspective. Our economic models show that: (1) Due to coordination difficulty, being profitable is not sufficient to guarantee a new architecture to be widely deployed; (2) A superior architecture may lose to another competing architecture which requires less coordination. Our model explains why IPv4 is hard to be replaced, why IPv6, DiffServ, CDN have different deployment difficulties, and why we observe the “Internet flattening phenomenon”. Our model also quantify the importance of incremental deployment mechanisms for the deployment of new Internet architectures. For architectures like DiffServ with which incremental deployment mechanisms are not available, people may consider a centralized mechanism to help the deployment. The designers of new architectures like NDN and XIA can also use our model to evaluate and improve the deployability of their design. Our model predicts that the current design of NDN and XIA are difficult to deploy, and MPQUIC will win over MPTCP.

Acknowledgment: The work of John C.S. Lui was supported in part by the GRF 14208816 and the Huawei Research Grant.

REFERENCES

- [1] Google, “Ipv6 stat.” www.google.com/intl/en/ipv6/statistics.html, 2018.
- [2] “Diffserv—the scalable end-to-end qos model,” Cisco, Tech. Rep., 2005.
- [3] Cisco, “Cisco report of cdn,” www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.html, 2017.
- [4] NDN, “Named data networking project.” named-data.net.
- [5] XIA, “expressive internet architecture project,” www.cs.cmu.edu/~xia.
- [6] MobilityFirst, “Mobilityfirst future internet architecture project.” mobilityfirst.winlab.rutgers.edu.
- [7] Li Ye, Hong Xie, and John C.S. Lui, “Quantifying deployability & evolvability of future internet architectures via economic models.” <https://arxiv.org/abs/1808.01787>, Tech. Rep., 2018.
- [8] L. S. Shapley, “A value for n-person games,” *Contributions to the Theory of Games*, vol. 2, no. 28, pp. 307–317, 1953.
- [9] F. Gul, “Bargaining foundations of shapley value,” *Econometrica: Journal of the Econometric Society*, pp. 81–95, 1989.
- [10] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, “Internet economics: The use of shapley value for isp settlement,” *IEEE/ACM TON*, vol. 18, no. 3, pp. 775–787, 2010.
- [11] X. Deng and C. H. Papadimitriou, “On the complexity of cooperative solution concepts,” *Mathematics of Operations Research*, vol. 19, no. 2, pp. 257–266, 1994.
- [12] D. Monderer and L. S. Shapley, “Potential games,” *Games and economic behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [13] D. M. Topkis, *Supermodularity and complementarity*. Princeton university press, 2011.
- [14] D. McFadden, “Conditional logit analysis of qualitative choice behavior,” *Frontiers in Econometrics*, 1973.
- [15] C. Alós-Ferrer and N. Netzer, “The logit-response dynamics,” *Games and Economic Behavior*, vol. 68, no. 2, pp. 413–427, 2010.
- [16] D. M. Frankel, S. Morris, and A. Pauzner, “Equilibrium selection in global games with strategic complementarities,” *Journal of Economic Theory*, vol. 108, no. 1, pp. 1–44, 2003.
- [17] F. Heinemann, R. Nagel, and P. Ockenfels, “The theory of global games on test: experimental analysis of coordination games with public and private information,” *Econometrica*, vol. 72, no. 5, pp. 1583–1599, 2004.
- [18] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, “The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse?” in *PAM ’08*. Springer-Verlag, 2008.
- [19] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan, “Are we one hop away from a better internet?” in *IMC ’15*. ACM, 2015.
- [20] M. K. Mukerjee, D. Han, S. Seshan, and P. Steenkiste, “Understanding tradeoffs in incremental deployment of new network architectures,” in *CoNEXT ’13*. ACM, 2013.
- [21] D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste, “XIA: Efficient support for evolvable internetworking,” in *NSDI ’12*. USENIX, 2012.
- [22] Microsoft, “Price of azure cdn,” azure.microsoft.com/en-us/pricing/details/cdn/, 2018.
- [23] Amazon, “Price of amazon shield,” aws.amazon.com/shield/pricing/.
- [24] M. Kühne, “Update on as path lengths over time,” labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time, 2012.
- [25] X. Jiang and J. Bi, “ncdn: Cdn enhanced with ndn,” in *Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2014.
- [26] IETF, “Tcp extensions for multipath operation with multiple addresses,” tools.ietf.org/html/rfc6824.
- [27] O. Bonaventure, “Measuring the adoption of mptcp is not so simple,” blog.multipath-tcp.org/blog/html/2015/10/27/adoption.html, 2015.
- [28] Q. De Coninck and O. Bonaventure, “Multipath quic: Design and evaluation,” in *CoNEXT ’17*. ACM, 2017.
- [29] A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kouranov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tenneti, R. Shade, R. Hamilton, V. Vasiliev, W.-T. Chang, and Z. Shi, “The quic transport protocol: Design and internet-scale deployment,” in *SIGCOMM ’17*. ACM, 2017.
- [30] G. Heal and H. Kunreuther, “Supermodularity and tipping,” National Bureau of Economic Research, Tech. Rep., 2006.
- [31] ISOC, “Ipv6 stat.” www.internetsociety.org/deploy360/ipv6/statistics/.
- [32] —, “World ipv6 launch,” www.worldipv6launch.org/.
- [33] —, “Six years since world launch, ipv6 now dominant internet protocol for many,” www.internetsociety.org/news/press-releases/2018/six-years-since-world-launch-ipv6-now-dominant-internet-protocol-for-many.
- [34] A. Bollapragada, “Ipv6 adoption in india,” www.nephos6.com/wp-content/uploads/2015/07/Study-of-IPv6-Adoption-in-India.pdf.
- [35] Xinhua, “China to speed up ipv6-based internet development,” english.gov.cn/policies/latest_releases/2017/11/26/content_281475955112300.htm, 2017.
- [36] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, “Providing public intradomain traffic matrices to the research community (<https://totem.info.ucl.ac.be/dataset.html>),” *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 83–86, 2006.
- [37] Caida, “The ipv4 routed /24 topology dataset.” www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml.
- [38] M. Roughan, “Simplifying the synthesis of internet traffic matrices,” *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 93–96, 2005.
- [39] A. Clauset, M. E. Newman, and C. Moore, “Finding community structure in very large networks,” *Physical review E*, vol. 70, no. 6, 2004.
- [40] S. Shenker, “Fundamental design issues for the future internet,” *IEEE JSAC*, vol. 13, no. 7, pp. 1176–1188, 1995.
- [41] R. E. Gilligan and E. Nordmark, “Basic transition mechanisms for ipv6 hosts and routers,” 2005.
- [42] R. Despres, “Ipv6 rapid deployment on ipv4 infrastructures (6rd),” 2010.
- [43] V. Kirilin and S. Gorinsky, “A protocol-ignorance perspective on incremental deployability of routing protocols,” in *IFIP Networking*, 2018.
- [44] T. Wolf, J. Griffioen, K. L. Calvert, R. Dutta, G. N. Rouskas, I. Baldin, and A. Nagurny, “Choicenet: Toward an economy plane for the internet,” *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, 2014.
- [45] S. Ratnasamy, S. Shenker, and S. McCanne, “Towards an evolvable internet architecture,” in *SIGCOMM ’05*. ACM, 2005.
- [46] H. Chan, D. Dash, A. Perrig, and H. Zhang, “Modeling adoptability of secure bgp protocol,” in *SIGCOMM ’06*. ACM, 2006.
- [47] P. Gill, M. Schapira, and S. Goldberg, “Let the market drive deployment: A strategy for transitioning to bgp security,” in *SIGCOMM ’11*. ACM, 2011.
- [48] S. Goldberg and Z. Liu, “The diffusion of networking technologies,” in *SODA ’13*. SIAM, 2013.
- [49] R. Cooper and A. John, “Coordinating coordination failures in keynesian models,” *The Quarterly Journal of Economics*, vol. 103, no. 3, 1988.
- [50] J. B. Van Huyck, R. C. Battalio, and R. O. Beil, “Tacit coordination games, strategic uncertainty, and coordination failure,” *The American Economic Review*, vol. 80, no. 1, pp. 234–248, 1990.
- [51] S. Morris and H. S. Shin, “Global games: Theory and applications,” *Advanced in Economics and Econometrics*, vol. 1, no. 67, 2001.