ChoiceNet: Network Innovation Through Choice

(Invited Paper)

George N. Rouskas[1], Ilia Baldine[2], Ken Calvert[3], Rudra Dutta[1], Jim Griffioen[3], Anna Nagurney[4], Tilman Wolf[5]
[1]Department of Computer Science, NC State University, Raleigh, NC, USA
[2]Renaissance Computing Institute, UNC Chapel Hill, Chapel Hill, NC, USA
[3]Department of Computer Science, University of Kentucky, Lexington, KY, USA
[4]Department of Finance and Operations Mgmt, [5]Department of ECE, University of Massachusetts, Amherst, MA USA

***Abstract*—Computer networks, in particular the Internet, represent essential infrastructure for business, government, military, and personal communication. Several recent trends in technology and network use have pushed the capabilities required of the Internet beyond what can be provided by the currently deployed infrastructure. To address these limitations, the network community has developed a variety of technologies to adapt the functionality of network protocols and services. A critical question that remains unanswered is how to integrate these technologies into an ecosystem that involves users, service providers, and developers in such a way that new ideas can be deployed and used in practice. In this work, we discuss how to design a network architecture where choices at different layers of the protocol stack are explicitly exposed to users. Our ChoiceNet system is based on three tightly coupled principles in that it aims to (1) encourage alternatives to allow users to choose among a range of services, (2) let users vote with their wallets to reward superior and innovative services, and (3) provide the mechanisms to stay informed on available alternatives and their performance. This approach ensures that innovative technical solutions can be deployed and rewarded, which is essential to encourage wide deployment of this architecture. Overall, our work does not aim at reinventing technical solutions to networking problems, but at developing a comprehensive system where these solutions can be incorporated and compete to allow the network to adapt to current and future challenges.**

I. Introduction

The Internet has been amazingly robust to changes in both technology and usage over the past three decades. A large reason for this success has been the fact that users are able to create and deploy a wide range of applications, devices, and services on end-systems around the edge of the network. In addition to enabling the Internet to support new uses that were not envisioned by its creators, this has also led to a vibrant economic environment for innovation at the application layer. The many successful Internet companies that have arisen in the last two decades testify to the ability of a platform on which novel ideas can be developed, tested, and adopted to lead to value creation.

Ironically, the same paradigm (i.e., the ability to create new functionality and let users choose winners and losers) is not presently supported inside the network. It is widely agreed that the current Internet architecture inadequately supports adaptation in the data and control planes [1]. In the early days of the Internet, innovations were deployed in the core by consensus among a small community of researchers and operators. With today's dramatically larger community, consensus is more difficult to achieve and innovation is for the most part limited to the edge. This limitation inhibits the development and deployment of new networking services, protocols, security designs, management frameworks, and other components that are essential to support the increasingly diverse systems, applications, and communication paradigms of the next-generation Internet.

Recently, the network research community, funded through programs including, in the US, NSF FIND and FIA [5], has developed new ideas and technologies for internetworking, including new functionalities for the control and data planes. For a comprehensive survey and comparison of research projects in future Internet architectures, the reader is referred to [10]. Software defined networking (SDN) [6] is another emerging network architecture that decouples the control and data planes. For the most part, however, these architectures focus on networking techology, not on economic interactions. For instance, while network control in the SDN architecture is programmable, the architecture itself does not provide any mechanisms to introduce competition and market forces in the control plane. A key open question, therefore, is how to bring together technical innovation inside the network by deploying meaningful mechanisms that allow network users (or their applications) to access these new features; and how to incorporate economics to ensure that innovation materializes by having someone pay for the new features.

Market forces have had a drastic effect on the shape of services and applications at the edge of the network. At the same time, there has been a wealth of studies that explore various economic issues that arise in the broad networking context, and, more specifically, the Internet [9]. Most of these studies attempt to analyze and understand the economic effects of existing networking technology. Our work has a different goal, namely, to integrate economic processes and interactions in the network architecture so that market forces can play out within the network itself. Accordingly, we propose a transformative shift in the design of networks that enables sustained innovation in the *core* of the network using economic principles. We believe that supporting *choice* is the key aspect of a network architecture that can adapt to emerging solutions

This material is based upon work supported by the National Science Foundation under Grants Nos. 1111040, 1111088, 1111256, and 1111276.

for current and future challenges [13]. Choice implies that users can select from alternatives that can be deployed dynamically into the network and reward those that address their needs. We use this interdependency between technological alternatives and economic incentives to create a competitive marketplace for innovative solutions that address current and future challenges in networking. Our work aims at the design, development, and prototyping of aspects of a next-generation network architecture where such choices and competition drive innovation at all layers of the protocol stack.

In this paper, we describe the outline of a network architecture that can provide choice as a core principle. The paper is organized as follows. In Section II, we discuss the three main principles underlying our design philosophy. In Section III, we describe the broad features of a network architecture that supports choice, and elaborate on the entities and their interactions within the architecture. We list use scenarios, including as they apply to optical networks, in Section IV, and we conclude the paper in Section V.

II. Choice As An Architectural Principle

Choice implies that the entities using the network can select from a range of alternative services that may differ in functionality, performance, and cost. Choices can and should appear at different layers in the protocol stack of a network, ranging from different communication paths to different protocols and application-layer services.

Choice can only arise if the network architecture supports the technology necessary for dynamic introduction of new alternatives. It is also necessary to put in place suitable economic processes to ensure that incentives trigger innovation and users can "vote with their wallets." Therefore, our system for a market-driven competition of networking functionality is based on the following three key principles we first described in [13]:

- **Principle 1: Encourage Alternatives.** The underlying network infrastructure must provide the building blocks to create different types of services and to create alternative services of the same type. Support for alternatives allows users to select the service that best meets their needs, and provides the best performance for their application. Implicit in this principle is the idea that, in contrast to the present Internet, where competition exists only at the application layer (if at all), users must not be "stuck" when the service they receive is not consistent with their expectations. Rather, they must be able to choose a different service provider, to better meet their expectations.
- **Principle 2: Vote With Your Wallet.** The underlying network infrastructure must provide building blocks so that users can financially encourage providers that offer superior (and often innovative) services, while discouraging providers that offer inferior services and fail to innovate. In other words, the "money protocols" needed for users to vote with their wallet (i.e., pay for good services) must be designed into the network. We believe that incentives and competition are crucial for the long-term health of the network: winning ideas will succeed, grow, and promote additional competition, while losing ideas fade away.
- **Principle 3: Know What Happened.** Distinguishing services and providers that perform well from those that do not is crucial to enabling robust competition. In a complex system like the Internet, determining what happened (i.e., whom to blame) when an end-to-end service does not meet user expectations can be a challenging proposition because providers may be operating at different layers and in several locations along an end-to-end path. The network must provide building blocks that allow users and providers to determine, and exchange information about, the performance they experience. Such "introspection" capability of the network also enables innovative network management and monitoring tools, which themselves can evolve over time.

These three principles interact in a cyclic process that results in increased competition, a faster pace of innovation, and better information for consumers, making it possible for users to make more informed decisions. We expect this "virtuous cycle" to be repeated over and over, on much smaller time scales than is possible in today's Internet, where customers are locked into their local provider and rely on the latter for everything. Furthermore, we note that omitting any element of the cycle destroys the effectiveness of the others. For example, if there are no other alternatives, there is no need to vote with one's wallet, nor does it help to know that the current service is not a good one. If there is no way to reward innovation and quality, i.e., to vote with one's wallet, providers will not have much reason to create new alternatives and improve service. If there is no way to "know what happened" (i.e., know what network services were or were not provided, or how they performed), it may be difficult or impossible to know which provider caused the problem, in turn making it difficult to vote with one's wallet. In fact, selecting among alternative services becomes a difficult process if no information about their performance is available (e.g., consider shopping online without any idea of the reliability or trustworthiness of the vendor).

III. The ChoiceNet Network Architecture

The sharp contrast between ChoiceNet's three guiding principles and the current Internet's lack of choice inside the network, call for an entirely new network architecture. Figure 1 illustrates schematically how the principles interact within a network, which we dub *ChoiceNet*, and the new features they support. ChoiceNet can be thought of as "network architecture" in that it requires a redesign of the data and control planes, as illustrated in the above figure and as we discuss shortly. However, it is not a complete architecture; instead, our aim is that ChoiceNet will fit with, and augment, existing and future architectures.

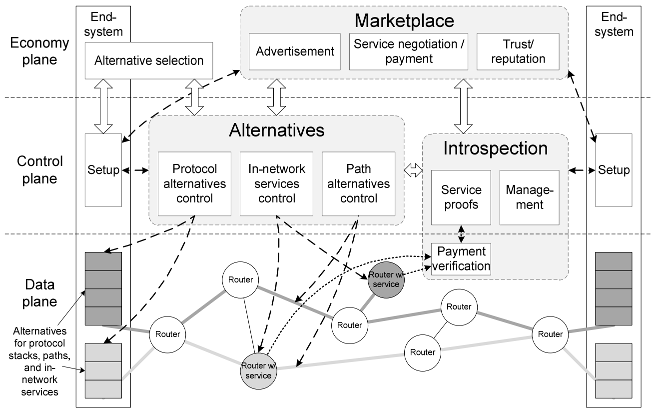

Fig. 1. Overview of the ChoiceNet architecture

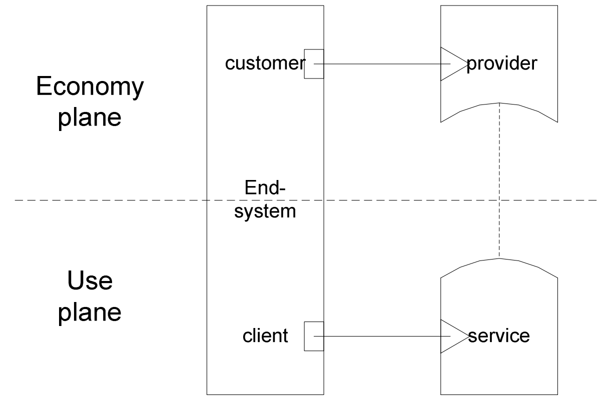

Fig. 2. Interfaces in ChoiceNet

A. The Economy Plane, Services and Entities

A key element of ChoiceNet is the new *economy plane* shown in Figure 1. The purpose of the economy plane is to expose choices throughout the protocol stack, and to enable and support economic transactions and business relationships (a "network services economy") over a wide range of time scales. Currently, business relationships take place out-of-band (i.e., outside the network itself) and typically operate at long time scales; for instance, customers select an Internet service, a particular data plan, or a music or video subscription and maintain the relationship with the respective provider for months or years. This state of affairs limits competition and has contributed to the "ossification" of the core protocols of the Internet stack [1]. The economy plane is responsible for supporting the types of incentives that drive innovation and change in the real world by making the process of establishing business relationships transparent, dynamic, and operable at short time scales. The economy plane supports advertisement of choices to users, as well as selection from among a set of alternatives. Economy plane protocols enable users to negotiate desired levels of service, and pay (compensate) specific providers for services. It also includes mechanisms to help establish the identity, level of trust, and reputation of the parties involved in a business relationship.

In ChoiceNet, we distinguish between technical and economic choices. For instance, consider the simple example of a movie streaming service. Technical choices involved in delivering such a service may include connection types, transport services, network paths, in-network caching, etc. Service providers may package technical choices and sell them as "experiences." In this case, the economic choices that an end user may face include, e.g., paying more or less for a particular movie experience. The economy plane is where technical and economic choices interact with each other to create a dynamic ecosystem of network services.

The concept of a "service" is a fundamental one in ChoiceNet, and is used to denote any functionality that can be realized within the network (i.e., "everything is a service"). Services may range from basic ones (e.g., representing simple network resources including bandwidth along a link, processing or storage capacity at a node) to more sophisticated ones that offer complex functionality (e.g., an end-to-end path with specific properties in terms of delay, jitter, etc). Services may be located in the data plane, control plane, economy plane, or a combination thereof. Services provide a benefit and have a cost; they can be created, sold, and verified. Importantly, a key aspect of the architecture is that multiple services may be combined together to create more advanced/complex services. Within ChoiceNet, we consider service composition as just another service. Service composition directly supports the first principle of ChoiceNet in that it allows the creation of additional alternatives (e.g., a suite of "movie experiences" that offer additional value to users) from a given set of basic services.

The marketplace is an important component of the economy plane in that it automates the process of offering and selecting services, thus supporting automatic service instantiation. Specifically, services are offered and advertised in the marketplace, and entities that require a particular service may query the marketplace to obtain a list of alternatives and competing offers. Hence, the marketplace is where economic transactions (including competition) take place. The marketplace may itself be viewed as a service, and, although this feature is not shown in Figure 1, ChoiceNet allows for multiple distinct marketplaces to co-exist so as to support implementation diversity.

A complex network such as the Internet must allow multiple entities to exist and interact. In ChoiceNet, we use a simple abstraction, depicted in Figure 2 that makes it possible to represent a wide range of entities and to construct complex economic relationships among these entities. Specifically, there are two interfaces along which various entities interact: a customer/provider interface in the economy plane, and a client/service interface in the data and control planes (collectively referred to as the "use" plane). In the economy plane, an entity acts as either a provider (offering

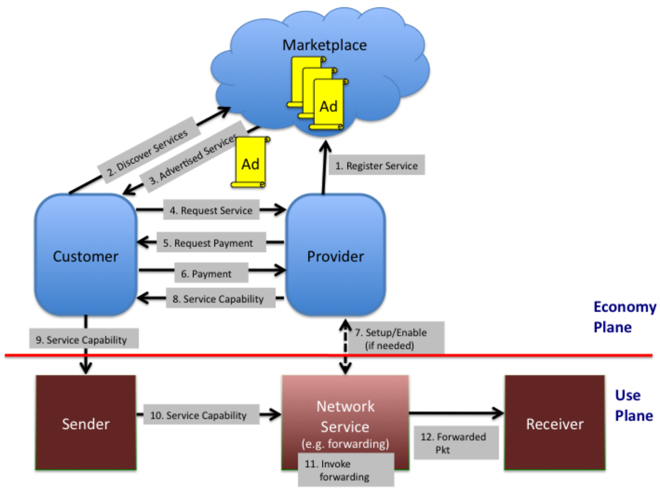

Fig. 3. Entities and Interactions in ChoiceNet

a service) or a customer (using a service); the role of each entity in the economy plane is determined by the established relationship (e.g., contract). In the use plane, entities interact as users (clients) of a service or providers of that service. We emphasize, however, that the client/service interface does not imply a client/server paradigm; indeed, the interactions may employ any appropriate communication paradigm.

Despite its simplicity, this abstraction allows more complex and realistic relationships to be constructed. For instance, it is possible for an entity ("integrator") to construct and offer a complex service from multiple simple services. An integrator acts in both the economy and use planes. In the economy plane, the integrator acts as the service provider for the customer of the complex service; at the same time, it acts as the customer of service providers from which it receives the simple services that are integrated into the complex service. These relationships are reflected in the use plane.

The separation of the two interfaces (in the economy and use planes) makes it possible to support entities that participate only in the economy plane. For instance, an entity may act as a "reseller" of services without implementing any network service itself. In this case, the participation of the reseller in the economy plane may be adding value, e.g., by finding the best service within parameters specified by the customer; in the use plane, the customer interacts as a client directly with the provider that actually implements the service.

The interactions among various entities in ChoiceNet within the economy and use planes are illustrated in Figure 3. When an end-system ("sender" in the figure) wishes to set up a connection with another end-system ("receiver"), it first queries the marketplace to obtain a list of services (and associated prices) that meet its specifications. After making a selection, the end-system contacts the service provider to request and pay for the selected service. More generally, at this point of the interaction, a contract is established between the customer and provider that outlines not only the payment but also the responsibilities of the service provider (e.g., in terms of QoS) in delivering the service. Once it receives payment (or a contract with a promise to pay has been established), the service provider sets up the service and provides credentials to the customer. All these interactions between customer and provider take place within the economy plane, as shown in the figure. The provider also activates the service in the use plane (e.g., by updating the state of routers) such that the end-system may start using the service. The credentials represent the right to access the service, and may take the form of, say, tokens that are included within the data packets for this communication. Note that variations of this basic interaction are possible, e.g., if an end-system already knows about the service and the provider, then it may not need to query the marketplace. Also, these interactions may be recursive depending on the business relationships we discussed earlier; for instance, when an end-system obtains service from an integrator, additional rounds of communication between the integrator and the service providers involved is necessary.

B. Mechanisms and Technology

We now describe some of the mechanisms, and relevant technology, that may be used to implement the ChoiceNet architecture. This discussion leverages and builds upon our past NSF FIND and GENI work including Postmodern Routing and Forwarding (Pomo) [2], data plane network services [12], and SILO [3].

1) Alternatives and Service Composition: As a starting point, one needs support for alternatives, including the ability to describe alternatives and combine them to create higher value services. A key goal of the ChoiceNet project is to develop building block services that, unlike the current Internet, support the offering of alternative network routes, in-network processing, and transport services. Users, that is, applications using the network will have some degree of control over the path(s) followed by their data through the network and how it is processed along the way. Following [3], to describe each service (alternative), the following information is provided: (1) the input/output interfaces, (2) the requirements on input data, and (3) the transformations performed on the input data. Using this information, users and their applications can assemble configurations vertically (i.e., protocol stacks) and horizontally (i.e., paths and in-network services), to suite their needs. To enable end-to-end composition of alternatives, we will extend the approach we have developed within the SILO framework [3]. Specifically, we have used ontology tools developed by the semantic web community to represent the semantics of services (i.e., the information identified above), along with composability constraints that represent the relationships among services. An ontology can be augmented with new information as new services and requirements arise, and can be used to perform semantic checks during composition.

2) Incenting Deployment of Alternatives: The existence of new techniques and protocols is not sufficient to ensure deployment; incentives are required to motivate providers to offer new alternatives. In today's Internet, customers compensate their local ISP to obtain a (basic, one-size-fits-all) transport service to all Internet destinations. Because there is no ability to compensate other providers directly, there is little reason for these providers to offer innovative services (except to their local customers). Any next generation Internet must include ways for users to compensate providers for services without regard to topology. More specifically, crucial building-block functions include those supporting the creation and maintenance of ongoing business relationships, and the verification of services rendered. ChoiceNet equips the network with two types of building blocks: (1) for service and payment authorization, and (2) for verification of payment. The key idea is to separate these two mechanisms, allowing them to occur at different time scales and in different locations. Our approach is based on capabilities, and leverages our past work on capability-based communication models [2], as well as related systems. These mechanisms are based on a simple idea: the provider issues to the customer a secret token, which is valid for some specified time interval. (The time interval can be viewed as a measure of the level of trust between customer and provider; the level of trust, in turn, may be based on the amount of compensation exchanged.) The customer attaches to each request for service a token, which is created using that secret and is bound to the service request via cryptographic hash. The provider can verify that the request came from the customer by verifying the hash computation and checking the time. The challenge is to scale this simple method up to a network with billions of users and millions of services.

3) Service and Payment Authorization: The goal of this mechanism is to enable users and providers to establish authenticated relationships and arrange for the exchange of compensation. To achieve "any-customer-to-any-provider" economic relationships that scale, a motivation delegation infrastructure can be used by users to "purchase" motivation tokens from providers. Motivation tokens can be thought of as capabilities that are checked by the provider before performing a service. We need a way for providers to give out motivation to the users of their service in a secure and scalable way. Consider the problem of incenting providers to forward users' packets. The straightforward application of the approach puts a token in each packet, which is verified by each router before forwarding the packet. Unfortunately, this requires each router to share a secret with each user. The introduction of brokers or "middlemen" can reduce the amount of storage required on both sides; further scalability is achieved by hierarchical delegation. To this end, we will extend the motivation mechanism in Pomo [2] to create a motivation delegation distribution tree.

4) Payment Verification: Although the policy, pricing, and negotiation decisions will occur outside the data plane, the services themselves must ultimately be performed as data traverses the network. Because the verification of "payment" must be performed in the data plane at "line speed' by all routers, we need a lightweight scalable mechanism. As described above, all the motivation tokens for a particular service originate from a secret key known to the service. Using that key, a service can generate a matching key for the particular branch of the delegation tree and verify the motivation token presented by the user. It can also verify that the motivation token is being used in the way it was authorized to use the service. Because of the way in which motivation tokens are generated, a service can pre-generate portions of the delegation tree allowing for fast motivation token lookups [2]. As noted earlier, we plan to build on the motivation delegation mechanism used in Pomo [2]. In Pomo, the basic service is a relay service that forwards packets from an ingress channel to an egress channel. Thus, verification occurs on a per-packet/per-relay basis, with each packet carrying its own motivation token. While the packet serves as the basic unit of verification for the Pomo forwarding service, other units of authentication may be better suited for other services. We will extend the Pomo approach to support other types of in-network services (storage, transcoding) and protocol services. The question is how to tie the motivation token to the service request. For some services, it may be carried out-of-band during a separate signaling phase used to invoke/establish the service. Verification may operate on packets, on application level framing units, or at a higher level (end-to-end session).

5) Introspection: Knowing What Happened: Having alternatives, and the mechanisms and economic power to choose between them, do not in themselves allow the user to make good choices that reward successful innovation, and bring about long term improvement. It is critical that the choice be informed by knowledge about what parts of the many services that determine the users communication experience have performed well on behalf of the user, and which have not. The problem is a significant one within ChoiceNet, because the user-network interaction is not a simple two-party interaction. The user, in undertaking any representative activity over the network, deals with several, perhaps many, service providers. Within ChoiceNet, customers may verify that they received the service that they paid for by having independent third-party providers offer two types of services: (1) Measurement service, and/or (2) Measurement analysis service. A measurement service uses active or passive measurements at various locations to collect data related to a service under test, and requires substantial investment in infrastructure. An analysis service, on the other hand, uses the data from the measurement service to derive the performance aspects of the service under test.

IV. USE SCENARIOS

To illustrate the ability of ChoiceNet to support innovation in the network, we briefly describe three use scenarios, highlighting the benefits of choice as a fundamental feature of the architecture. The scenarios also show the different scales at which our three principles apply, thus illustrating the generality of innovation through choice.

A. Scenario 1: Network-Level Choices

Network virtualization is one of the key technologies for supporting multiple different networks on the same infrastructure [11]. While the technology for such a solution is maturing, it remains an open question how best to close the loop between customers, service providers, and the underlying infrastructure. A network architecture that is based on ChoiceNet principles would solve this problem: Customers (e.g., large companies or industry sectors) could specify their needs for a custom virtual network, service providers could compete to offer different solutions, and the business goes to the provider (or several) that provides the best service. For example, many existing and emerging classes of high-end applications involve complex, intensive computations on large data sets in a manner that requires coordination of resources residing at several geographically dispersed sites [8]. To enable these applications, lightpaths along end-to-end paths across multi-domain networks must come up and go down, based on user requirements, and over short time scales. Using the economy plane, different service providers can offer their networks and added-value services (e.g., various levels of path diversity, restorability, various layer-2 technologies, etc.); using the feedback from the network, users can monitor the quality and performance of their service provider and choose (and thus financially reward) the one that fits them best. Thus, *ChoiceNet provides mechanisms for the deployment of networks with specific characteristics.*

B. Scenario 2: Connection-Level Choices

The levels of customization of individual network connections is rather coarse in the current Internet; essentially one can choose between TCP and roll-your-own over UDP. Techniques for customizing protocol stacks and in-network services have been studied and, in some cases, deployed [3], [12]. To make these technologies successful (thus leading to an innovative and improved network), users need to be able to make informed choices among offerings, and providers need to be able to benefit from users selecting their alternatives. For example, different network providers could offer customized connections for high-volume data transfers over high-speed optical links. Their offerings may differ in transport protocol [7], connection performance (bandwidth, delay, loss rate, etc.), in network services (multicast), and in cost. Users (or their applications) may select among these offerings and switch providers in case a provider cannot deliver the promised functionality or performance. In the existing network, such competition does not exist, is too coarse-grained, or is tied to users choosing different end-system applications in case they are unsatisfied. Thus, *ChoiceNet enables competition between providers within the framework of the network.*

C. Scenario 3: Service Virtualization

In existing networks, there is often a clear separation between users and network service providers. As the Internet grows more diverse, this distinction becomes increasingly blurred. Users may have access to resources that they can share (e.g., temporarily unused wireless spectrum, storage for content distribution, etc.) and thus effectively become providers at a very small scale [4]. A forward-looking, riskier goal of our research is to employ ChoiceNet principles to support such recursive service offerings and ensure that economic incentives lead to innovation at this level. Thus, *ChoiceNet enables a paradigm shift toward a network where everyone can be a provider.*

V. Concluding Remarks

We have presented, ChoiceNet, an architecture that augments existing networks to enable choice and the establishment of dynamic business relationships. We argue that the core components of ChoiceNet, including a new economy plane, will lead to innovation and foster an ecosystem of network services. We expect that the architecture we described will lead to a different Internet. Specifically, instead of being dominated by a few large and vertically integrated network providers, we believe that an economy plane-enabled Internet will see an overall increase in the number of providers, including new ones that find niche applications in which to compete and be successful. Such a development would democratize the Internet and ignite competition and innovation into the core of the network.

References

[1] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. *IEEE Computer*, 38(4):34–41, April 2005.

[2] K. Calvert, J. Griffioen, and L. Poutievsky. Separating routing and forwarding: A clean-slate network layer design. In *Proceedings of Broadnets 2007 Conference*, September 2007.

[3] R. Dutta, G. N. Rouskas, I. Baldine, A. Bragg, and D. Stevenson. The SILO architecture for services integration, control, and optimization for the future internet. In *Proceedings of IEEE ICC*, pages 1899–1904, June 2007.

[4] N. Feamster, L. Gao, and J. Rexford. How to lease the Internet in your spare time. *SIGCOMM Comput. Commun. Rev.*, 37(1):61–64, 2007.

[5] D. Fisher. US National Science Foundation and the Future Internet Design. *ACM Computer Communication Review*, 37(3):85–87, July 2007.

[6] Open Networking Foundation. Software-defined networking: The new norm for networks. https://www.opennetworking.org/images/stories/downloads/white-papers/wp%-sdn-newnorm.pdf, April 2012.

[7] Data Transport Research Group. Survey of transport protocols other than standard grid TCP. Global Grid Forum, 2005.

[8] G. Karmous-Edwards. Global E-Science collaboration. *IEEE Computing and Scientific Engineering*, March/April 2004.

[9] L. W. McKnight and J. P. Bailey (Eds.). *Internet Economics*. MIT Press, 1997.

[10] J. Pan, S. Paul, and R. Jain. A survey of the research on future internet architectures. *IEEE Communications Magazine*, 49(7):26–36, July 2011.

[11] A. Wang, M. Iyer, R. Dutta, G. N. Rouskas, and I. Baldine. Network virtalization: Technologies, perspectives and frontiers. *IEEE/OSA Journal of Lightwave Technology*, 31(4):523–537, February 2013.

[12] T. Wolf. In-network services for customization in next-generation networks. *IEEE Network*, 24(4):6–12, July 2010.

[13] T. Wolf, J. Griffioen, K. L. Calvert, R. Dutta, G. N. Rouskas, I. Baldine, and A. Nagurney. Choice as a principle in network architecture. In *Proceedings of ACM Annual Conference of the Special Interest Group on Data Communication (SIGCOMM)*, pages 105–106, Helsinki, Finland, August 2012. (Poster).