# ChoiceNet: Toward an Economy Plane for the Internet*

Tilman Wolf
ECE Department
University of Massachusetts
Amherst, MA, USA
wolf@ecs.umass.edu

Jim Griffioen, Ken Calvert
CS Department
University of Kentucky
Lexington, KY, USA
{griff,calvert}@netlab.uky.edu

Rudra Dutta, George Rouskas
CS Department
North Carolina State Univ.
Raleigh, NC, USA
{dutta,rouskas}@csc.ncsu.edu

Ilya Baldin
RENCI
University of North Carolina
Chapel Hill, NC, USA
ibaldin@renci.org

Anna Nagurney
Operations and Inform. Mgnt.
University of Massachusetts
Amherst, MA, USA
nagurney@isenberg.umass.edu

## ABSTRACT

The Internet has been a key enabling technology for many new distributed applications and services. However, the deployment of new protocols and services in the Internet infrastructure itself has been sluggish, especially where economic incentives for network providers are unclear. In our work, we seek to develop an "economy plane" for the Internet that enables network providers to offer new network-based services (QoS, storage, etc.) for sale to customers. The explicit connection between economic relationships and network services across various time scales enables users to select among service alternatives. The resulting competition among network service providers will lead to overall better technological solutions and more competitive prices. In this paper, we present the architectural aspects of our ChoiceNet economy plane as well as some of the technological problems that need to be addressed in a practical deployment.

## 1. INTRODUCTION

The network architecture of the Internet has proven to be sufficiently flexible to enable an amazing array of distributed applications and services. However, with expanding requirements for the Internet, including security, privacy, support for the Internet-of-things, etc., new network architectures are being explored. As noted in the seminal work by Clark et al. [13], a critical consideration for any new network architecture design is the need to explicitly allow for real-world tussles to take place within the architecture. One such tussle revolves around the economic relationships between entities operating and using the Internet.

While the networking research community continues to develop many novel network protocols and services, there have been considerable challenges in actually deploying innovative solutions in the Internet, specifically in-network services, such as path selection across domains, secure routing/naming, etc. As an example, consider the very limited deployment and use of multicast, which generally does not fit well in the Internet's general one-size-fits-all/all-you-can-eat business model [14]. Therefore, it is essential that new network architectures pay attention to economic interactions that are likely to become tussles, and enable them to play out in "protocol space."

The current Internet reflects certain economic relationships, for example between network service providers. However, these relationships are at best implicitly encoded in BGP routing policies [10]. This approach has led to complexity and slow convergence of global routing protocols [18]. Moreover, it poses a high barrier of entry for entities that do not participate in routing to be part of the economy surrounding in-network services.

In our work, we aim for an *explicit* representation of economic relationships among entities in the Internet. (We consider the *network* ecosystem—the provisioning of communication services using shared infrastructure—as distinct from the application-level ecosystem, which simply uses the network as a communication channel.) We propose an *economy plane* that enables anyone to offer *network services* and to set up *contracts* for their use. The main aspects of this economy plane are one or more *marketplaces* where services are offered and sold; a method for describing network services such that they can be compared and composed; and the mechanisms necessary to implement and enforce contracts at various time scales.

One fundamental change that we target as an outcome from the use of an economy plane is that users (or their applications) can select from a number of different network offerings rather than being limited to whatever a single provider offers them. This "refactoring" of services to enable greater *choice* should promote competition among providers for price and quality; this is known to lead to benefits for consumers and the economy [28].

## 2. CHOICENET ARCHITECTURE

The main goal of ChoiceNet is to provide technologies and protocols to support economy plane transactions and relationships in the Internet. Our primary focus is on traditional network layer features (e.g., routing, quality), but we aim for general mechanisms that can be used for emerging functions as well (e.g., in-network storage, stream transformation, location-based services). The idea is to use such economic mechanisms to reward services that are perceived as useful by customers. We first discuss the goals and architectural fixed points of ChoiceNet in this section and then describe implementation details in Section 3.

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

## 2.1 Motivation and Vision

The goal of ChoiceNet is to enable choices and the associated economic relationships between entities in the network. ChoiceNet makes it possible for network service providers to compete for customers and be rewarded for quality and innovation. In today's network, money enters the network ecosystem only around the edges: consumers (individuals or enterprises) pay *access* providers for Internet service. Most access providers, in turn, pay other ISPs to carry their traffic to/from the rest of the Internet. Indeed, most end-to-end traffic in the Internet traverses at least three distinct service providers.

Thus, in today's Internet (i) no single provider controls all end-to-end paths; (ii) money flow between providers is outside the architecture and by necessity changes slowly; and (iii) traffic flow is constrained *at the granularity of providers* to follow the money flow. The result is that transit providers have neither means nor incentive to compete via new service offerings, and consumers have essentially no control over the service they receive or its quality. A central thesis of ChoiceNet is that enabling money flow to follow traffic flow (instead of vice versa), coupled with greater support for choice among end-users, should lead to increased provider competition and more innovation.

ChoiceNet's economy plane aims to give assurances to providers that they can compete for customers and be compensated for the services they render. At the same time, ChoiceNet provides users with the ability to select from a set of offerings and combine them to form complex services, thereby separating services that are currently entangled in the current Internet. Key to such an architecture is the ability to market services and then form or dissolve business relationships on (potentially small) time scales. Moreover, ChoiceNet must enable providers (and consumers) to prove (or verify) that the contracted service was rendered as promised.

To illustrate this vision with a specific example, we briefly discuss how choice can motivate the deployment of innovative new services and protocols for streaming video content. In the current Internet, there are a number of different content providers that offer video streaming services (e.g., Netflix, Hulu, Amazon, etc.). While users can choose a video content provider, users cannot choose how the network handles the streaming traffic produced by the provider as it traverses the network. As a result, users can only hope that they receive the content with enough quality of service to have an enjoyable experience. Content providers also have little or no control over the network when sending their content and instead focus on coding and quality adaptation techniques to adjust to changing network conditions.

In a ChoiceNet-enabled Internet, we envision that users will explicitly select network services from a marketplace of competing service offerings. Every aspect of the network can be marketed and sold as a service. Moreover, it is possible to combine services, such as specific network paths, in-network storage, packet forwarding prioritization, etc., into a tailored movie-watching package. These packages can be crafted by knowledgeable users, their applications, or, more likely, by service providers. Users can choose among different packages, experience their quality, and decide which package they want to continue using and paying for. We describe the operation of ChoiceNet in the context of this example in more detail in Section 3.5.

While this scenario requires significant changes in how economic relationships are established between customers and providers, there is also need for technological change. Services need to be general, so that they can be combined into useful packages. Also, the network infrastructure itself needs to be diverse enough to offer a variety of choices at the network layer. While the Internet offers a wide range of end-system services, the set of end-to-end network/transport layer services available is rather small. However, recent advances in software defined networking, programmable networks, and cloud computing make it possible to offer such alternative services.

Not only should ChoiceNet support alternatives for consumers to choose from and purchase, but it also needs to support a variety of economic relationships. Providers may bundle and resell services offered by others, adding value in the process; in doing so they act as both customers and providers. For example, today's mobile virtual network operators (MVNOs) provide cellular network access to users through short-term service contracts, but do not operate their own infrastructure. Instead they resell network access from other providers. Similarly, crowdroaming enables network access to participants by sharing access points. These current-day examples, however, lack a general architecture (or in some cases specific economic models), which is what ChoiceNet aims to provide.

## 2.2 ChoiceNet Components

Here we describe the key components of ChoiceNet in more detail.

### 2.2.1 Entities

ChoiceNet comprises the *economy plane* and the *use plane*. The former is where customers and providers interact to establish economic relationships for network services; the latter is where services are realized and corresponds to the traditional data plane and control plane.

Interactions that occur in ChoiceNet's two planes can be described simply as:

**Customer-provider relationships in economy plane:** *Customers* interact with *providers* to obtain access to one or more services. An entity may act as a client to one side and as a customer to the other side. Such transitivity enables the composition of more advanced services without the need for the provider to have access to physical infrastructure.

**Client-service relationships in use plane:** Providers enable *services* based on economy plane agreements so (authorized) *clients* can use them.

The interfaces and their relationships are shown in Figure 1(a). A service constructed from two other services is illustrated in Figure 1(b). It is also possible for services to be offered by economy plane entities that have no physical presence in the use plane and only act as resellers of services.

### 2.2.2 Services

ChoiceNet services range from simple bit pipes to payload processing functionality, from data transmission to data storage. These services are offered in marketplaces and can be obtained by anyone—end-system users or providers of service compositions.

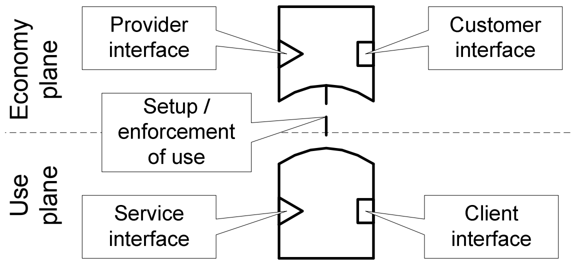


(a) Interfaces of entities in ChoiceNet.

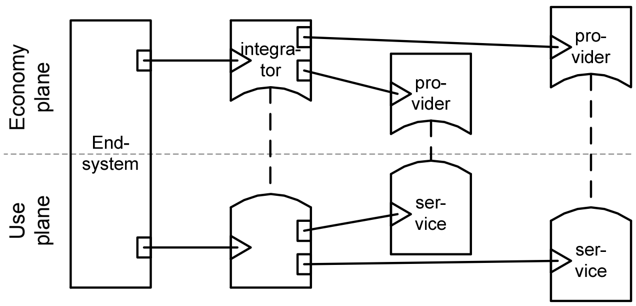


(b) Construction of service from multiple other services.

**Figure 1: Representation of entities in ChoiceNet.**

Providers start out with local services (e.g., local transit policies) and price them (e.g., using local conditions). ChoiceNet entities (e.g., "middlemen") then act in the economy and use plane to combine and bundle these service to end-to-end services (e.g., similar to combining pathlets into end-to-end connections [16]). The various end-to-end compositions compete with each other, as do local providers, which leads to the desired competition and innovation.

It is important that services be described in a reasonably generic fashion to enable comparison between services. This aspect is critical to support the goal of enabling competition among service providers.

Network services typically move data from one place to another, with data possibly modified in the process. Thus, a first-order service description in ChoiceNet consists of two parts:

- Service end-points: Service end-points determine the location of the input(s) and output(s) of a service. The abstraction level of the location description is not specified a priori. Depending on the type of service, location can range from an identifier, such as a network address, to higher-level semantics, such as an end-user identity. The only requirement is that the end-point can be named in some context. It is in the interest of the service provider to use naming(s) that can be understood by potential customers.

- Service semantics: The semantics of a service can be described by specifying the following two aspects of a service: (i) **Input requirements:** These requirements state what the service expects to receive on its inputs in order to function correctly. Examples range from requiring valid IP packets to requiring traffic with particular payloads (e.g., H.264-encoded video) and specific quality-of-service (QoS) parameters (e.g., 1 Mbps limit). (ii) **Transformations:** The service transformations specify how the service modifies the data received on the inputs. In the simplest case, traffic is moved from input to output without modification. In more complex scenarios, payloads may be inspected (e.g., for security) or modified (e.g., video transcoding). The service semantics can, but do not have to, specify QoS constraints imposed by a service.

More complex services (e.g., storage, multicast, etc.) can be represented by straightforward extensions. For example, a storage service consist of two components, one with an input and no output end-point (corresponding to "write") and one with no input and an output end-point (corresponding to "read"). A multicast service would have a single input and multiple output end-points.

Based on these service descriptions, services can then be composed into more complex service offerings. More complex compositions, such as video transcoding, in-network caching, etc., rely on semantic abstractions of services. Approaches for automated service selection and composition are discussed in Section 3.1.

### *2.2.3 Contracts*

Contracts are used in the economy plane to set up the economic exchanges that precede the set up and use of services in the use plane. Providers may ask for any of multiple different *considerations* in return for providing a service. We envision that many considerations correspond to real-world currencies, but they do not have to. An example of non-currency consideration is proof that customer has access rights to service because of organizational affiliation. The customer may choose which type of consideration to "pay" to the provider.

The exchange of consideration needs to be such that the provider is satisfied with receiving the consideration. This can be achieved through a third party (e.g., credit card processing company, Paypal, etc.) that is able to ensure that accounts were debited/credited appropriately, but can also be done in other ways.

Enforcement of contracts is critical to ensure the viability of the economy plane. This enforcement occurs in two dimensions: (1) providers want to ensure that use plane traffic has matching economy plane contracts (i.e., "has been paid for") and (2) customers want to ensure that they receive the service they paid for (i.e., "know what happened"). We discuss how contracts are verified in the use plane in Sections 3.3 and 3.4.

### *2.2.4 Marketplace*

To choose, customers must first know what choices are available for their network services. The purpose of the marketplace component of the ChoiceNet architecture is to act as a "service commons," a meeting ground for provider advertisements and user requirements. Each provider of service advertises each service they offer. Note that in the marketplace concept, as in other things, ChoiceNet attempts to reflect and automate real world economic interactions—in this case, shopping. However, in the real world such shopping depends largely on previously established business relationships. Automation allows software agents acting on behalf of the real user's preferences to search the marketplace, making it possible for a large number of providers' advertisements to be considered by each of a large body of users, thus flattening the marketplace. This in turn allows innovative alternative providers to enter the marketplace even in the absence of an established relationship, and build new ecosystems of support.

It is necessary for the user to know which of these services can be used as alternatives for each other, and which may be used in conjunction with each other; thus ChoiceNet defines common minimal semantics for advertisements in the marketplace. Each advertisement has a set of attributes. Core attributes, for which each advertisement must provide values (from among possible values listed in a shared ChoiceNet dictionary), indicate the type of service as well as details like provider name, certificates, advertised cost, a link or handle to access the provider, etc. Based on the specific type of service, additional attributes provide information on what other services can or must be used in conjunction. The user (or their software agent) can use attribute values of advertisements to plan their desired network service (see Section 3.1).

The ChoiceNet architecture provides marketplaces as places in which providers and users meet and supplies the minimally necessary semantics for them to exchange information about the services. The negotiation of the consideration for services is left to providers and users. Auctions and other market mechanisms can be implemented as value-add policies by providers that act as aggregators of multiple marketplaces. The ChoiceNet architecture views marketplaces as providers, thus permitting hierarchical arrangements of marketplaces offering service bundling and auction services. We note that this creates a tension with service delivery transparency, which can be resolved by higher level providers communicating the details of services acquired from the lower level providers to the users. Such transparency will present a market edge to those providers who offer it.

## 2.3 Identities, Trust, Security

Customers and providers must be able to identify each other in both economy and use planes in order to transact business and access/provide services. A global namespace is therefore a necessary component of the ChoiceNet system. A flat namespace with distributed assignment and so-called "self-certifying" associated public/private key pairs has several attractive features (e.g., it makes authentication of communications easy), but requires a trusted resolution service to map more human-friendly specifications of such IDs. In keeping with the goal of universal applicability of ChoiceNet, it is possible to envision a set of trusted *identity providers* who provide unique, certified IDs, but this brings issues of cross-service trust and interoperability: the cost of supporting multiple formats grows quadratically in the number of such providers. We note that ChoiceNet has this problem in common with other FIA projects (e.g., MobilityFirst [2], XIA [1]), and therefore remains agnostic (so far) regarding most aspects of namespaces.

It is worth noting that each additional level of naming and resolution implies an additional layer of *trust* that must be bootstrapped and reckoned with. ChoiceNet offers the advantage of making the extent of trust explicit: the customer's compensation, and the mechanism through which it is transferred to the provider, are both exposed to both parties. The extent to which the customer's trust in the provider is justified can be revealed by the verification service.

## 2.4 Economic Principles

The economy plane in ChoiceNet enables and supports economic transactions and business relationships in a "network services economy" over a wide range of time scales. In order to assess different business models, it is imperative to construct and evaluate rigorous computable mathematical game theory models that capture the interactions of the various economic agents, that is, the providers (content, transport, etc.) and the consumers, over space and time.

Toward that end, we have been utilizing variational inequality theory [20] and projected dynamical systems theory (see [23]) to quantify the competition among the various providers, as they compete in prices, quality, and quantity, which may differ, depending upon the tier of provider (see, e.g., [21, 22] for specific illustrations). The resulting competition is of benefit for consumers and for the economy [28] since the competitive market transforms self-interest into a force for public good and competition maximizes productivity and social welfare through the optimal allocation of capital and labor in the economy. Since these seminal insights in [28], economists agree that competition leads to the lowest price and the highest quality products and provides consumers with a greater spectrum of choices [27].

We note that, although the focus of this paper is on ChoiceNet and bilateral economic relationships as between customers and providers, through composition, novel services in which there are multilateral relationships (see, e.g., [11, 12]), with associated pricing and contracts are also feasible.

# 3. CHOICENET OPERATION

To illustrate how the general principles of ChoiceNet can be translated into concrete network operation, we discuss the steps required for communication in this section. The interaction of components in ChoiceNet is shown in Figure 2. There are three major actions that need to be performed to set up services in the economy and use plane: *planning*, *provisioning*, and *usage*.

## 3.1 Planning

After providers advertise their services in the marketplace, the first step in setting up communication in ChoiceNet is the planning step. During planning, customers explore the available choices for their communication needs. While customers consider the cost and quality of different choices, no contracts are set up and no networking resources are committed. For planning to work, providers first have to advertise their available services in one or more marketplaces.

Wrapping marketplace advertisements in Web Service Definition Language (WSDL) is one possibility of an adequate existing mechanism that facilitates expressing which services can be composed after which others, allowing such planning. Service details not involved with composition need not have a WSDL representation, but can be encapsulated inside. As mentioned above, values for such attributes come from a shared ChoiceNet dictionary, in effect creating a ChoiceNet service description language (similar to NDL-OWL models used in ORCA [8] and GENI [25], or their conceptual predecessors).

The main algorithmic problem in planning is to find one or more combinations of available services that meet the requirements of the customer. Similar to the service specification discussed in Section 2.2.2, requirements specify the ser-

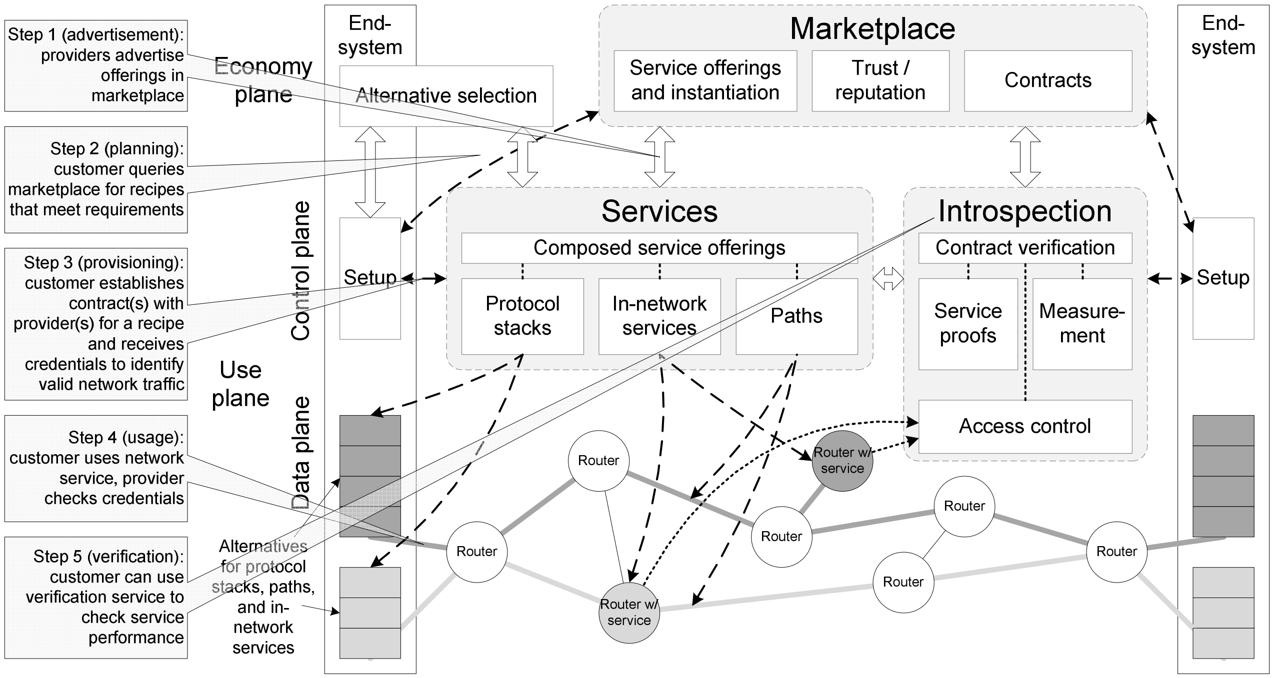


**Figure 2: ChoiceNet architecture and operation.**

vice end-points as well as semantics and QoS requirements. The planning algorithm, which can be implemented by the marketplace or by the customer, then searches through available services to find a composed "recipe" that meets these requirements. The output is then a set of Pareto-optimal service combinations, or recipes. An example implementation of a service planning algorithm can be found in [15].

## 3.2 Provisioning

During provisioning, consideration is exchanged in the economy plane and resources are committed accordingly. Orchestration and provisioning actions take place across multiple providers, setting up dependencies between them and allowing each individual provider to perform its own provisioning actions. Since service recipes can be represented by a directed acyclic graph, the creation of complex orchestrated arrangements of paths does not require multi-step negotiations between adjacent peer providers, simplifying the setup. The output of the provisioning step is the set of authorization tokens that enable the customer to prove policy-compliance in the use plane to get access to the service.

In an example scenario for provisioning of path segments, provider domains can *produce*, *consume* and *translate* labels and an external entity must orchestrate the process of constructing a path, by letting e.g., *label producer* domains allocate and provision their labeled paths first, passing the labels on to *label consumers* for stitching to them. At its simplest, e.g. as used in GENI today [9], it boils down to determining common VLAN tags for layer 2 paths, where a transit provider may be a label producer, creating a path with one or two distinct labels at each end of the path, passing this information to edge cloud providers to attach their compute resources to those tags to create a complete arrangement of connected resources.

## 3.3 Usage

The previous two steps, planning and provisioning, involve the economy plane and the marketing/purchasing of services. While the economy plane is needed to create choice and competition, the user is ultimately interested in *using* services. The use plane is simply the collection of mechanisms that enable services to verify that a requested usage has been authorized in the economy plane. Note that the use plane may not need to know a particular user's identity or what a particular user is allowed to do. (Indeed, the whole point is to confine the decisions based on policy to the economy plane, and specifically the provisioning step.) Consequently, the use plane does not need to deal with the issues of user management.

Provisioning produces a set of *tokens* that enable the user to convince the service that it is entitled to access/use the service. The form of this proof of policy-compliance may vary with the security and performance requirements, e.g., from a simple cleartext identifier to a cryptographically secure proof of possession of a secret or capability [17, 26, 30, 32]. The proof is included with each request for service, which could be as frequent as every packet (e.g., for a packet forwarding service) or as infrequent as the first message in a session (e.g., for an in-network file storage service), and may be "bound" to the request to ensure tokens are not stolen or misused.

Upon receiving a service request, the service verifies that the token authorizes the current request, and if successful, performs the requested operation (which might involve forwarding the request onward to the next stop in a composed/stitched set of services).

## 3.4 Introspection and Verification

The availability of alternatives and the economic power to choose between them does not inherently bring about good choices that reward successful innovation. It is critical that users have the means to evaluate their choices.

When a user or software agent plans the user's desired network service (such as watching streaming video) the expectations of the user are often defined in emergent terms, such as some quantification of user experience ("the video should never freeze for more than a tenth of a second, and there should be no drastic loss of resolution in any frame"). Thus, the marketplace advertisement must include appropriate attributes to specify the performance that the customer can expect from the service. From the customer's perspective, the only potential for enforcement of such a promise lies in their ability to take their business elsewhere, if they are not satisfied by the service provided by a provider, i.e., "vote with their wallet." To do this effectively, the customer must be able to decide whether they received the service they had been led to expect; in other words, the performance attributes advertised must be *independently verifiable*.

The typical customer is able to make measurements only on an end-to-end basis. This is not sufficient to separate the performance of the different component services. One straightforward approach is for a third-party measurement service provider to maintain Points of Presence at Internet Exchange Points where different component service providers hand over user traffic from one to the other. Note that, in view of the ChoiceNet paradigm, it is unnecessary for a customer to know the performance of every router or every link along the path, as so much existing literature has focused on. Measurements only need be provided at points when one provider hands over the responsibility of transporting (or otherwise processing) the customer's data to another provider. A sequence of even a few such measurements can enable the customer to draw useful conclusions regarding "whom to blame" when end-to-end performance turns out to be not as promised, as in the example from [7] that we mention below.

Neither the idea of measurements to monitor service quality, nor the idea of third-party measurement providers, is new. The passive and active measurements research community has long investigated innovative monitoring techniques, and initiatives such as MLab (originally Google MLab) [5] have attempted to create a consortium exchanging measurement data gathered in the network in an open format, to empower users. Much earlier attempts such as that of CAIDA [29] may also be seen in this light. However, the adoption and penetration of such methods in application or platform (OS) software has not progressed well, because measurement at flow granularity is expensive, and measurement is often done not for profit.

This barrier to realization is likely to be lowered, if not removed, if such service is perceived to be of value to customers who are willing to pay for it. Flexible measurement transport architectures that allow a tunable or demand-based measurement transfer, such as perfSONAR, or IPFIX (the successor of NetFlow), already exist, and can be adopted for the purpose. In [7], we have shown how such an approach can allow a customer to analyze simple timestamp information for a subset of the packets in the flow generated at each interface between successive component service providers, and perform the difficult emergent task of jitter apportionment—deciding which provider introduced a critical amount of jitter causing a movie to freeze.

## 3.5 Example Operation

Revisiting our movie streaming example from Section 2.1, we briefly discuss how the various ChoiceNet operations would occur in this specific context.

First, network providers advertise their network transit services in the marketplace. The service description of the network transit service would specify the end points (e.g., router interfaces) and the service semantics (e.g., QoS parameters).

When a user wants to connect to a video provider to stream content, their application needs to establish an end-to-end connection. To identify the choice of available and matching service compositions, the application queries one of the available marketplaces by providing the requirements for the service recipe (i.e., one end-point located at the user's computer, other end-point located at content provider's server, minimum QoS requirements). The marketplace then performs search and computation operations to identify valid service compositions that meet the user's request. For example, there could be two different paths to the video server, one with higher guaranteed bandwidth but also higher cost to the user. A set of these service offerings are sent to the user's application, where the user, the application, their operating system, or an administrator pick a suitable service offering (e.g., the higher-bandwidth path).

Once a service offering has been chosen, the marketplace is informed of that choice and sets up contracts between the customer and the service providers (possibly acting as a trusted intermediary). For example, users and marketplace could use the APIs provided by Paypal or credit card companies to exchange funds. After the contract is in place, the providers set up their routers to enable the service that was purchased (e.g., by adding suitable rules into their SDN switches to make a certain path available). The content provider then transmits the video stream and the user receives it with the quality that they paid for.

If the user desires to be in a position to know whom to blame in case the video quality is unacceptably bad, the user would search for, choose, purchase and provision a verification service in the same manner as for the path services or content streaming services. Verification service providers would advertise in the marketplace just like other providers and specify details of what quantities they are capable of verifying at what network locations. The user would supply the composed data service specifications to verification service provider, which would (after purchase) begin monitoring the data flow, and provide occasional or on-demand performance reports of the path providers and content providers who are providing the video service to the user.

Clearly, there can be many variations on this example, all of which are enabled by the ChoiceNet architecture: there can be multiple marketplaces; the content provider pays for a better network path and charges the user indirectly; customer and providers can set up contracts directly without using a marketplace as trusted intermediary; etc. It is important to note that the ChoiceNet architecture enables all these alternative realizations of the architecture without prescribing a specific one.

# 4. DISCUSSION

Here, we attempt to address some questions commonly raised about ChoiceNet.

### *How will ChoiceNet improve things?.*

Going back to the streaming example of Section 2.1, today both the consumer and the content provider have very few "knobs" they can manipulate to control their satisfaction. The consumer may choose a different source for the content; The streaming service can innovate with codecs and protocols, but is at the mercy of access and transit providers, who do not share in its (advertising-derived or subscription-based) profits. Indeed, increased success for the streaming service correlates with more traffic and increased *cost* for the ISPs. (Thus, we have "net neutrality" tussles playing out in "policy space," rather than within the architecture.)

We believe this is a consequence of the fact that compensation for network services enters the ecosystem *only* in the very place where consumers have the fewest choices, and barriers to entry are highest, viz., at the access provider. The ChoiceNet hypothesis is that enabling compensation to flow to services more or less independent of location will encourage competition and innovation. This is a complex multi-dimensional proposition, for which we do not expect to provide a rigorous proof. Besides the many technical challenges, there are likely even more non-technical hurdles.

For the video streaming example, it may or may not be the case that building dedicated infrastructure to replicate content across all access networks is "better" in many dimensions than paying a service to find and compose longer paths from available bandwidth in real time. It seems clear, however, that innovation in network services can only be hampered in the long run by the current highly constrained money flow.

### *How will ChoiceNet scale?.*

Scalability of ChoiceNet will depend on the timescales at which choices are made, the spectrum of options supported, how choices are made, and the threat models to be protected against. While the intent is to impose as few constraints as possible and admit a full range of solutions, we do not expect that the full range will be realized in practice. For example, we do not expect that all (or even most) users will need or want to choose most aspects of their network services most of the time. Rather, the point of admitting the full spectrum is to constrain innovation as little as possible, and to allow the "sweet spot" of the flexibility-efficiency tradeoff to adapt with technology changes.

It is clearly unrealistic to expect humans to be "in the loop" for choices related to individual packets, or even for short-lived flows. In general, the intent is for choices to be either (i) automated and controlled via human-specified policies, which can be changed as desired, or (ii) presented in a form similar to what is experienced today, i.e., which link to click on. The former admits preconfigured policies that default to familiar options (best-effort, access-provider-selected transit, for example).

Clearly, the extra flow of information and consideration inherent in the ChoiceNet model represent overhead that will ultimately constrain the achievable rate and granularity of choice. Moreover, because incentives are the whole point, additional overhead will be required to ensure trust, at least in some cases. However, real-world examples of systems that conduct negotiations on very short timescales are not unknown. For example, online ad exchanges already handle billions of transactions per day, auctioning ad space on a timescale of a hundred milliseconds or so. Recent work has shown that such auctions can be made cryptographically auditable at modest cost [6].

Business relationships present another scaling challenge. It is not reasonable to expect even a few dozen transit providers to deal individually with billions of potential consumers, nor will consumers wish to deal with billing by multiple transit providers—not to mention other types of providers, as network-based services expand to include storage and compute cycles. ChoiceNet allows for third parties to aggregate network services on one side and access to consumers on the other, thus providing "retail" classes of service as is common in other kinds of markets. Indeed, in some cases, third parties with a global view are required to hide the local, non-fungible nature of some services. For example, transit providers need only implement local transit policies (i.e., relaying packets from ingress links to egress links, regardless of origin or destination—as long as they are compensated). However, bandwidth is strictly a local resource, so *path services* will be needed to aggregate relays with available capacity into end-to-end paths for resale to end-users.

### *What about market clearing in ChoiceNet?.*

Market clearing is a key issue for ChoiceNet. Ideally there would be a single, trusted global clearinghouse. Although this would be the simplest and most efficient solution, it is unlikely. More likely would be a scenario in which a limited number of "clearing providers" provide some or all of clearing, risk management, and billing services. Obviously some such entities of this kind already play a large role in today's economy (Visa, PayPal, etc.) Recent interest in distributed proof-of-work-based transaction ledger protocols [24] raises the intriguing possibility of a decentralized accounting solution.

## 5. RELATED WORK

Research efforts addressing future Internet designs are extremely diverse in terms of both technical approach and geography. Due to space constraints, we only discuss sister projects funded by the NSF Future Internet Architecture (FIA) program, while acknowledging that there are many other ongoing efforts in the US, EU, China, Japan, and other countries.

The Named Data Networking (NDN) project [3] has content distribution as the key theme. Recognizing that the primary use of the Internet today is no longer for host-to-host communication but rather for content delivery, NDN focuses on *what* the users want (i.e., content) rather than *where* to find it (i.e., IP addresses). The MobilityFirst project [2] is motivated by the shift from fixed end-points (that the original Internet architecture was designed to connect) to the current demand for access to mobile devices and services. The main goal of the project is a new architecture for a pervasive system that will connect mobile devices and vehicles among each other and to the infrastructure services. The NEBULA project [4] takes an approach starting from the premise that storage and computation will continue migrating into the "cloud." Therefore, the main goal of the project is to design an architecture that enables a future core Internet to interconnect data centers that provide highly reliable and secure "computing utility" services. The eXpressive Internet Architecture (XIA) project [1] has as its objective the creation of a single architecture that supports secure and trustworthy communication among a diverse set of principals that include hosts, services, and content, as well as entities that may emerge in the future. XIA specifies intrinsic mechanisms, including secure identifiers, to establish trust and secure the communication among principals; hence defining the APIs and semantics of this communication is a major research thrust for the project. These network architecture provide the technical choices that can be represented in the economy plane of ChoiceNet. We therefore view our project as complementary to these ongoing efforts.

There has been other work on introducing economic mechanisms into the Internet. Yang [31] describes a scheme for billing on essentially a per-packet basis by letting users "bid" for bandwidth. However, that scheme relies on the money flow implied by the current hierarchy of transit providers, while ChoiceNet aims to accommodate arbitrary business relationships. Resource allocation based on economic principles has also been done in grid computing [19]. That work uses auction-based approaches for simple (i.e., not composed) computing services.

# 6. SUMMARY

Choice of network services is a critical requirement for future Internet architectures that aim to enable economic competition and innovation. In our work, we develop the architecture for an economy plane for the current and future Internet. ChoiceNet provides the principles and mechanisms to offer and buy network services in marketplaces using short-term contracts, including service description and contract enforcement.

The ChoiceNet economy plane is designed to work in conjunction with existing and emerging networking technologies and thus provides a general approach to explicit representation of economic relationships in the network architecture.

# 7. REFERENCES


[1] eXpressive Internet Architecture project. `http://www.cs.cmu.edu/~xia/`.

[2] MobilityFirst future internet architecture project. `http://mobilityfirst.winlab.rutgers.edu`.

[3] Named Data Networking project. `http://named-data.net`.

[4] NEBULA: Future internet architecture. `http://nebula-fia.org`.

[5] The virtual computing lab. http://www.measurementlab.net.

[6] Angel, S., and Walfish, M. Verifiable auctions for online ad exchanges. *SIGCOMM Comput. Commun. Rev. 43*, 4 (Aug. 2013), 195–206.

[7] Babaoglu, A., and Dutta, R. A verification service architecture for the future internet. In *Proc. of the 22nd IEEE International Conference on Computer Communications and Networks (ICCCN)* (Nassau, Bahamas, Aug. 2013).

[8] Bal, H. E., Kaashoek, M. F., and Tanenbaum, A. S. Orca: A language for parallel programming of distributed systems. *IEEE Trans. Softw. Eng. 18*, 3 (Mar. 1992), 190–205.

[9] Baldine, I., Xin, Y., Mandal, A., Heermann, C., Chase, J., Marupadi, V., Yumerefendi, A., and Irwin, D. Networked cloud orchestration: A GENI perspective. In *IEEE GLOBECOM Workshops* (2010), pp. 573–578.

[10] Caesar, M., and Rexford, J. BGP routing policies in ISP networks. *Network, IEEE 19*, 6 (2005), 5–11.

[11] Castro, I., and Gorinsky, S. T4P: Hybrid interconnection for cost reduction. In *INFOCOM Workshops* (2012), pp. 178–183.

[12] Castro, I., Stanojevic, R., and Gorinsky, S. Using tuangou to reduce IP transit costs. *IEEE/ACM Transactions on Networking*. To appear.

[13] Clark, D. D., Wroclawski, J., Sollins, K. R., and Braden, R. Tussle in cyberspace: defining tomorrow's internet. *SIGCOMM Comp. Comm. Rev. 32*, 4 (Aug. 2002), 347–356.

[14] Diot, C., Levine, B. N., Lyles, B., Kassem, H., and Balensiefen, D. Deployment issues for the IP multicast service and architecture. *Netwrk. Mag. of Global Internetwkg. 14*, 1 (Jan. 2000), 78–88.

[15] Dwaraki, A., and Wolf, T. Service instantiation in an Internet with choices. In *Proc. of the 22nd IEEE International Conference on Computer Communications and Networks (ICCCN)* (Nassau, Bahamas, Aug. 2013).

[16] Godfrey, P. B., Ganichev, I., Shenker, S., and Stoica, I. Pathlet routing. *SIGCOMM Computer Communication Review 39*, 4 (Aug. 2009), 111–122.

[17] Kambhampati, V., Papadopoulos, C., and Massey, D. A taxonomy of capabilities based DDoS defense architectures. In *The 9th IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)* (2011), pp. 157–164.

[18] Labovitz, C., Ahuja, A., Bose, A., and Jahanian, F. Delayed internet routing convergence. *SIGCOMM Comput. Commun. Rev. 30*, 4 (Aug. 2000), 175–187.

[19] Lai, K., Rasmusson, L., Adar, E., Zhang, L., and Huberman, B. A. Tycoon: An implementation of a distributed, market-based resource allocation system. *Multiagent Grid Syst. 1*, 3 (Aug. 2005), 169–182.

[20] Nagurney, A. *Network Economics: A Variational Inequality Approach*, second and revised ed. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1999.

[21] Nagurney, A., Li, D., Wolf, T., and Saberi, S. A network economic game theory model of a service-oriented internet with choices and quality competition. *NETNOMICS: Economic Research and Electronic Networking 14*, 1-2 (Nov. 2013), 1–25.

[22] Nagurney, A., and Wolf, T. A Cournot-Nash-Bertrand game theory model of a service-oriented internet with price and quality competition among network transport providers. *Computational Management Science*. To appear.

[23] Nagurney, A., and Zhang, D. *Projected Dynamical Systems and Variational Inequalities with Applications*. Kluwer Academic Publishers, Sept. 1996.

[24] Nakamoto, S. *Bitcoin: A peer-to-peer electronic cash system*. `http://bitcoin.org/bitcoin.pdf`, 2008.

[25] National Science Foundation. *Global Environment for Network Innovation*. `http://www.geni.net/`.

[26] Raghavan, B., Verkaik, P., and Snoeren, A. C. Secure and policy-compliant source routing. *IEEE/ACM Trans. Netw. 17*, 3 (2009), 764–777.

[27] Shapiro, C., and Varian, H. R. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Review Press, Boston, MA, 1999.

[28] Smith, A. *An Inquiry Into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell, London, UK, 1776.

[29] WWW-document. Cooperative Association for Internet Data Analysis. *http://www.caida.org/*.

[30] Yaar, A., Perrig, A., and Song, D. SIFF: A Stateless Internet Flow Filter to Mitigate DDoS Flooding Attacks. In *IEEE Symposium on Security and Privacy* (May 2004), pp. 130–143.

[31] Yang, X. Auction, but don't block. In *Proc. of NetEcon '08* (Seattle, Washington, August 2008).

[32] Yang, X., Wetherall, D., and Anderson, T. A DoS-limiting network architecture. *SIGCOMM Comp. Comm. Rev. 35*, 4 (2005), 241–252.