Converging Choice and Service in Future Commodity Optical Networks using Traffic Grooming

Rudra Dutta*, George Rouskas*, Ilia Baldine†
*Department of Computer Science, North Carolina State University, Raleigh, NC, USA
†Rennaissance Computing Institute, University of North Carolina, Chapel Hill, NC
Email: rdutta@ncsu.edu, rouskas@ncsu.edu, ibaldin@renci.org

Abstract

The problem of providing an agile, energy-aware, flexible optical network architecture is one of the challenges in optical networking in the coming decade. A key element in this challenge is the balancing of the benefits to customer and provider, and creating an agile system capable of reflecting both provider and customer interests on an ongoing basis as network conditions change. In this paper, we articulate how the traditional optical networking research area of *traffic grooming* may be combined with recent advances in Internet architecture, specifically a proposed Future Internet architecture called *ChoiceNet*, and empowered by the recently emerged concept of *software defined networking*, to make some key contributions to this problem.

I. Introduction

Optical transport networks have formed the lowest level of the core of the planetary communication networks, whether voice or data, over the last decades, and can be expected to do so indefinitely into the future. Access networks and local networks continue to use electronic equipment suitable for complex computation in the path of data flows, and connect to the core optical networks using opto-electronic devices. The explosive growth of mobile networking has focused attention on wireless networks, while at the same time increasing and underscoring the dependence of planetary communications on high-bandwidth optical cores.

The changing landscape of planetary networking has posed several emerging needs for the optical backbone that are new. The optical transport must support an increasingly larger range of bandwidth needs, over more diverse timescales. As dynamic provisioning algorithms operate on heuristic principles, over time the succession of traffic demand arrival and departure can create suboptimalities in network resource utilization; thus the network must re-optimize itself from time to time. As timescales decrease, this can happen over short intervals such as hours, and the re-optimization must become an integral part of network operation, rather than a separate network management function. Finally, to allow a diverse and stable ecosystem of network operators and service providers, different niches for collaborating and competing businesses must exist, such that different business entities with different value proposition, business models, risk tolerance etc. can leverage each others' offerings to present a rich set of offerings to the customer.

In this paper, we present recent research threads that can combine to produce an elegant coherent vision of the future optical service network. Below we describe these areas briefly, and indicate how they fit together in our vision.

II. Traffic Grooming as Mapping

The art and science of converging available technologies, electronic and optical, for the access and core, for network-wide benefit, has been known as *traffic grooming* in the optical networking research literature. This area generated a good deal of research in the past, which largely focused on optimizing the number of wavelengths, or OEO conversions, in keeping with the techno-commercial need of the time. This has given rise to the narrow definition of traffic grooming as the act of multiplexing sub-wavelength flows into wavelength channels. The literature on grooming has been surveyed variously [1]–[4], and aspects of the research topic reviewed comprehensively [5].

The range of literature on grooming contained in the surveys and tutorials above shows it to be a broad area, focused on multi-layer approaches to traffic engineering and resource placement/optimization, characterized by explicit representations of constraints and opportunities specific to optical layer technologies. More recent work that falls under the umbrella of grooming includes so-called "green Grooming" that attempts to consolidate or distribute traffic over the network with an eye to reducing energy expenditure, either at individual points or network-wide. The disparate need for cooling system power required for differing choices that use different tradeoffs for electronic and optical technology provides an example of the input to such considerations.

As optical technology evolves, new devices or techniques bring new opportunities and constraints, and grooming approaches must be revisited to address them, say advances in elastic wavelengths. On the other hand, the need of the users whose data flows constitute the access network traffic that flows through the core

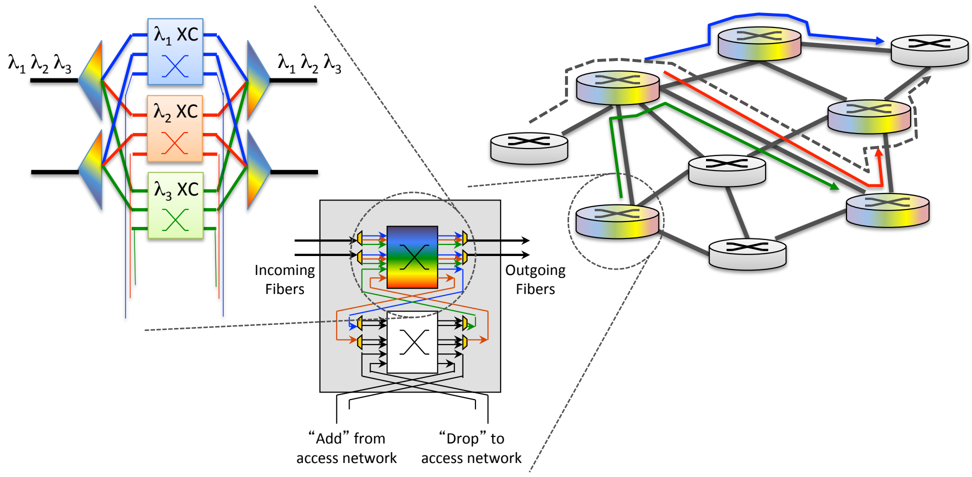

Fig. 1: Traffic grooming converges available technologies of core and access for network-wide mutual benefit

also evolve, causing changes in the traffic demand characteristic, that have to be represented in such problems – the growing need for connection mobility is such an emergent change, prompted by (i) the rise in the volume of data representing individual users traffic, and (ii) the high degree of mobility that is becoming the norm for individual users even as they access high-bandwidth channels. The optical network of the future must be agile, and able to react quickly, on an ongoing basis, to changing user demands; it must also continually reconfigure itself and decide what services to offer so that they can be profitably aggregated or engineered in the network.

Grooming techniques can thus be used to balance various optimization criteria. At a given time, the bulk of the Opto-Electro-Optic forwarding equipment available in a provider's infrastructure may be committed to carrying currently resident traffic flows, and it may be preferable for the provider to overprovision a new traffic request as an optically switched path. This is the oldest traffic grooming scenario. However, at a different time, bandwidth might be the precious commodity, and it might be more expeditious for that provider to advertise cheaper, low-bandwidth paths without strong bounds on delay variance, so as to leverage its packet-switching capability. The provider can run grooming scenarios based on such diverse parameters as its current available equipment, expected demands in the near future, profit margins for various classes of services, energy costs associated with each class, etc., and decide at each recomputation epoch the optimal set of services to offer.

Provided grooming algorithms can be used to decide the best service offerings, there needs to be a mechanism to allow the provider to realize such offerings in an agile manner, and for the customer to be able to find the right offering from the dynamically changing set of offerings. We discuss these next.

III. ChoiceNet - Service Marketplace

Even as the current Internet enables a range of services and distributed applications that grow ever broader and more vareigated, several limitations of its architecture have become apparent as billions of humans and devices are connected through it. One key challenge is the discrepancy between the mechanisms by which technology is deployed in the Internet and the business models surrounding these processes.

No valuable network function can be assumed to be free, or made available as a public service. In monetory transactions for services, competition improves the collective benefit. Competition is more effective when the marketplace is larger. A small marketplace can become a buyers market or sellers market more easily, and also because a larger variety of buyers means a larger variety of producers and sellers can be sustained. In particular, small producers, who in a marketplace of only a few large buyers will either become captive to one of them or be driven out, may flourish in a large variegated marketplace containing small buyers who can be directly sold to. These reflections prompt the goal of ChoiceNet to introduce architectrual entities into the Internet to enable fine-grain economic interactions.

We have previously considered these issues in the course of our ChoiceNet project [6]–[8], under the umbrella of the NSF Future Internet Design program [9]. Briefly, the ChoiceNet project proposes the introduction of architectural entities into the current architecture of the Internet to enable an "economyplane" that allows the presentation of competing offerings for various networking services, the formation of contracts for the various services that make up the entirety of a user's network needs, and tracking the performance of each provider in meeting their parts of the contracts. The workflow of service offering and consumption is represented by three principles of ChoiceNet: (i) "Enable Alternatives": provide the mechanisms and building blocks for users to be presented with, and to choose different services easily, (ii) "Know what happened": allow the user to find out which provider to blame if service expectations are not met, (iii) "Vote with your wallet": to provide the means for secure, scalable and fine-granularity payment protocols to allow the user to promote better performing providers. For ease of exposition, we can focus on the case when this “voting” takes the shape of money changing hands, mediated through automated protocols; but in general we speak of *consideration* instead, which may be cash in some cases but may be in kind in others: barter is an example, but so is a situation where the very utilization of the service offered by a provider represents value to the provider, due to accumultion of goodwill, leverage for advertising (the famous “supplying eyeballs” model of e-commerce), or beta-testing a product.

The ChoiceNet architecture provides a common platform for service providers to advertise their services and for customers to easily discover, negoitate and pay for these services. The component empowering the advertisements for services is called the “marketplace” where service providers register their services and customers discover them via querying the marketplace. Once a customer decides the service to purchase, further ChoiceNet interaction between customer and provider creates a contract, with the customer receiving a *token* of some form. These interactions constitute what we term the “Economy Plane” of ChoiceNet, and are all paralleled by real-world interactions that take place, but outside the network architecture, today. Subsequent to this, the customer may use the token to obtain the desired service in the same manner as would take place in the current Internet - we refer to this as the “Use Plane”. Such services could be for endpoint services, path or pathlet services, or even in-network processing services. However, in the ChoiceNet view, the customer would undertake individual contracts with each of the providers of such services along the complete path to compose the entire service, and would be financially rewarding each provider, thus creating incentive for innovation at all such service points, endpoints, paths, and intermediate processing/storage services. In reality, automated software acting on behalf of the user would undertake these decisions, translating high-level goals of user experience into low-level ones.

IV. Agile Optical SDN

The idea of Software-Defined Networking (SDN), where management and intelligence functions of network control is separated from the data plane, and resides on separate controllers that can be instantiated in software (therefore allowing policy to change at short timescales), has become popular recently. The OpenFlow [10] system has become generally accepted in the researcher community, and is gaining ground in the practitioner community, as an extant SDN system, though it is important to note that it does not represent all the different aspects or scope of SDNs. It is a limited function SDN that aims at carefully delineated functions it attempts, and others that it does not. Its focus is on centralizing routing policy, more correctly on centralizing the point of application of a new routing policy, by separating it from the mechanism of forwarding itself. This centralizes the routing policy, in the manner of path computation elements, rather than distibuting it over all participating forwarding engines. Thus changes in routing policy (even drastic ones such as switching from a simple shortest path to a complex traffic engineered approach) can be easily achieved at the central controller, and does not require reconfiguring (or worse, upgrading) every forwarding engine. Since the controller is not in the path of data flow, and does not have to operate at wirespeed, it is natural to implement it in software, so that it is cheaper and more agile. The forwarding engine maintains only a “flowtable” in which flow entries are indexed by a “match”, which are partial or complete tuples drawn from a union of IP, a few popular transports, and a popular link/frame standard. The match may be extended, starting with Version 1.2, to include other header quantities as TLV tuples [11]. This creates the possibility of using OpenFlow with such OpenFlow Extensible Match (OXM) tuples to control agile optical equipment to dynamically create services of varying granularity and isolation characteristics (lightpaths, time-slotted opto-electronic paths, statistically multiplexed paths), as determined by the grooming algorithms.

We note that the architecture we are envisioning for the ecosystem involves re-factoring of functionality that are analogous in multiple dimensions. SDN (or OpenFlow, in particular) provides separation of the control and data plane functions of the network, as described above. In an analogous fashion, ChoiceNet provides separation of service advertisement from service use, thus allowing grooming algorithms to address the problem of service definition not inline with the operation of the use plane interactions of the purchase and provisioning of such service. We describe our complete vision next.

V. The Converged System

The above, in combination, allows us to envision an ecosystem of mutual benefit. The primary entities in this ecosystem are the customer, the provider, and the marketplace. In this view, service re-sellers and bundlers are represented as a combination of customer and provider entities. These entities are similar though not precisely identical to the corresponding ChoiceNet entities described above, and their interactions are somewhat simplified versions of ChoiceNet, customized for this specific application space. Figure 2 shows these entities and their interactions. The customer and the provider are engaged in mutual value exchange; typically, the customer needs service that the provider has the bandwidth and switching infrastructure to produce (its stock-in-trade), and the customer provides some consideration, often cash, that the provider values. The marketplace is an entity that serves as the rendezvous between provider capabilities and customer needs. Since the rendezvous represents value to both customer and provider, the marketplace can be assumed to be a cooperatively realized non-profit entity, especiallly since it need be little more than an agile directory with search capabilities. The labeled solid arrows in Figure 2 show the different possible interactions between the entities, and the dotted arrows represent the most typical order of the various interactions in the life-cycle of the convergence of network capabilities and customer requirements, and the resulting collaborative dynamic optimization of the network resources. Although we have indicated an order, the interactions are obviously asynchronous and can take place in any order, and each is executed an indefinite number of times, on an ongoing basis, for a given provider or customer.

Fig. 2: Entities and interactions for convergence ecosystem

A straightforward manner in which the marketplace accomplishes the rendezvous between the provider and the customer is indicated by the sequence of interactions marked “list offerings” (by provider, to market-place), and “retrieve offerings” (by customer, from marketplace). This would naturally be followed by “make buying decision”, which represents the contract between customer and provider that consists of payment for a connection or bandwidth service for some quantum of time, and the provisioning of such service by the provider. The provider would naturally use this information to “update resource availability” internally, and “update listing/prices” in the marketplace. We do not show them, but there would naturally be a complementary interaction for the customer to release the contract, or signal end of service; this would also be followed by a re-assessment of available network resources by the provider.

While this is a natural sequence of interactions, it does not represent anything different from the business interaction that takes place today, although by means of human interaction rather than thorugh automated signaling. Our key observation is that at every epoch, a provider can use traffic grooming algorithms to decide the most optimal set of service offerings to list in the marketplace, in light of its remaining network resources (bandwidth, electrical and optical switching capability, buffers). Typically, the set of service offerings (granularity, guarantee of delay bound, jitter, etc.) for a provider have to be “well-known” in order for prospective customers to weigh different offerings from the same or different providers in light of their own requirements, and make a buying choice. However, this does not allow the provider to be very responsive to current network conditions; the marketplace mechanism allows the provider to dynamically and continuously customize its service offerings to best leverage network resources available at any given time. Further, providers typicaly analyze historical data on service usage to determine what service usage is likely to sell best. With an agile marketplace, it is possible for such a provider to hedge its bets by monitoring the uptake of an offered service in near real-time, and re-evaluate its resource provisioning strategy, or combine strategies, if necessary. The introduction of grooming algorithms makes it possible for a provider to not only be responsive to a changing set of resident traffic demands (hence available resources), but also to customize its offerings to optimize figures of merit such as equipment or bandwidth utilization, revenue generation, energy efficiency, etc. Thus the agility of the marketplace can be used to support dynamic re-optimization of the network. Finally, the network provider must be able to provide “handles” to the dynamically offered bandwidth services that conform to a few easily accessible standard technologies for the customer, such as MPLS labels, VLAN IDs, or even wavelengths, while internally mapping these, at differing times, to bandwidth tunnels of different granularities, different combination of forwarding technology (e.g. optical or electronic, various levels of electronic), and different buffering/scheduling strategies resulting in differing quality of experience (e.g. time-slotted or statistically multiplexed). To achieve this, the provider can utilize optical networking equipment that incorporates SDN mechanisms, either proprietary to the vendor of the equipment, or using an open platform such as OpenFlow with standardized extensions specific to optical equipment as we have suggested above. In the latter case, the openness of the control plane can be leveraged to re-use the same primitives for customer interactions also; in other words, the customer could simply federate its OF controller with that of the provider, and use any suitable “handle” for the service, leaving it to the federation between the OF controllers to correlate the marketplace (economy plane) interaction with the traffic flow (use plane).

One other set of interactions is shown in Figure 2. This enables a customer “express interest” in a potential service type even when it is not currently offered in the marketplace by a given provider. This is similar to customer-drive product specification such as tendering a bid. Such a tender can also have associated information indicating what level of consideration the customer is willing to or reasonably expects to pay. A provider can retrieve such tenders, and take this into account when using grooming algorithms to compute the next set of service offerings to advertise, specifically relating them to the tender, so that a customer can at a later time retrieve offers specifically listed in response to the interest. This provides a further, explicit channel for the customer needs to affect the network configuration, and provide mutual benefit.

A. *The GENI-IMF Project*

In a rudimentary form, the Integrated Measurement Framework project executed by us in GENI [12] explores the integrated scenario we have just described. The GENI-IMF project is described in detail at [13]. Briefly, it involved optical layer measurements made at a GENI optical substrate, communicated to a stack protocol running inside a GENI slice, which dynamically exercised path choices as well as optical power choices to stabilize video quality in the face of wavering optical port power at an intermediate node. While this project pre-dates our thinking on ChoiceNet and thus does not use ChoiceNet standard interactions, and uses proprietary control technology for the dynamic control of optical equipment (rather than OpenFlow or other open standard), it contains a corresponding set of entities (GENI slice: customer; optical substrate: provider; power control options and path options: marketplace), and interactions (choose new path: make buying decision) under varying conditions of network resources (wavering port power). This provides us with confidence in our basic premise of an automated rendezvous of dynamic service offerings and service requirements for future optical networks.

VI. Conclusion

Optical backbones provide the bandwidth necessary for all planetary communications, but are often not perceived by the end consumer. As such, the economy, control signaling, and provisioning timescales, have all remained isolated and disconnected between backbone networks providing bulk bandwidth and commodity networks providing service to consumers. This has created barriers for the emergence of innovative and timely, consumer-responsive service offerings. In this paper, we have examined a possible concatenation of existing research ideas to form an ecosystem that allows providers and consumers to cooperatively enable efficient use of available network resources to mutual benefit. Interesting research problems present themselves in specifying such a system in detail, and designing it in realistic terms. We believe that such research will be undertaken by the optical networking research community in the near future, and will have a transformative effect on the study and practice of optical networking.

The authors would like to thank the other PIs of the ChoiceNet project, Drs. Tilman Wolf, Kenneth Calvert, Jim Griffoen, Anna Nagurney, and collaborator Dr. Shu Huang, for their valuable feedback. This work was supported by NSF NeTS award 1111276.

References

[1] R. Dutta and G.N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, Nov/Dec 2002.

[2] K. Zhu and B. Mukherjee. A review of traffic grooming in WDM optical networks: architectures and challenges. *Optical Networks Magazine. March April 2003; 4(2): 55-64*, 2003.

[3] E. Modiano and P.J. Lin. Traffic grooming in WDM networks. *IEEE Communications Magazine*, 39(7):124–129, Jul 2001.

[4] Shu Huang and Rudra Dutta. Dynamic traffic grooming: the changing role of traffic grooming. *IEEE Communications Surveys and Tutorials*, 9(1):32–49, 2007.

[5] R. Dutta, A.E. Kamal, and G.N. Rouskas. *Traffic grooming for optical networks: foundations, techniques and frontiers*. Springer, New York, 2008. Optical networks series.

[6] Tilman Wolf, James Griffioen, Kenneth L. Calvert, Rudra Dutta, George N. Rouskas, Ilia Baldine, and Anna Nagurney. Choice as a principle in network architecture. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, SIGCOMM ’12, pages 105–106, New York, NY, USA, 2012. ACM.

[7] George N. Rouskas, Ilia Baldine, Kenneth L. Calvert, Rudra Dutta, James Griffioen, Anna Nagurney, and Tilman Wolf. Choicenet: Network innovation through choice. In *Proceedings of ONDM 2013*.

[8] A.C. Babaoglu and R. Dutta. A verification service architecture for the future internet. In *Computer Communications and Networks, 2013. ICCCN 2013. Proceedings.22nd International Conference on*, 2013.

[9] Nsf future internet architecture project. *http://www.nets-fia.net/* .

[10] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. Openflow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.

[11] Open Networking Foundation. Openflow switch specification. https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.3.1.pdf.

[12] Geni: Exploring networks of the future. *http://www.geni.net/* .

[13] Geni integrated measurement framework. *http://groups.geni.net/geni/wiki/IMF* .