ABSTRACT

UDECHUKWU, ROBINSON N. Toward Realizing a Choice-based Co-Optimizable Networking Paradigm. (Under the direction of Dr. Rudra Dutta.)

We investigate the network global performance impacts of interplay of provider strategies for pricing and path selection, and customer strategies for accepting path alternatives, in a choice-based network scenario with hybrid optical packet and circuit switching capability. In preliminary empirical units of work, we demonstrate the feasibility of fine-grain traffic control with commercial equipment, and of marketplace based customer and provider interaction to manage a hybrid optical network. Then we investigate the hypothesis that such interaction can enable network co-optimization, in which collective customer satisfaction is improved at the same time as network global performance metrics. On the provider side, we propose a baseline pricing algorithm that only takes into account the cost of the network resources that would be required to realize a possible path alternative for any given customer request, and two other pricing algorithms (one of which is parameterizable) that each attempt to reflect the current occupancy state of those resources into the pricing. Also on the provider side, we propose different strategies for provisioning paths, pre-empting packet wavebands for the benefit of circuits as necessary. On the customer side, we define different utility-based customer strategies for deciding how to select from among available path alternatives. We undertake simulation studies to explore various scenarios of such interactions, including one in which competing providers cooperate to provide a richer set of alternatives to customers, and present our results and conclusions.

© Copyright 2021 by Robinson N. Udechukwu

All Rights Reserved

Toward Realizing a Choice-based Co-Optimizable
Networking Paradigm

by
Robinson N. Udechukwu

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Computer Science

Raleigh, North Carolina

2021

APPROVED BY:

Dr. George Rouskas

Dr. Mihail Sichitiu

Dr. David Thuento

Dr. Rudra Dutta
Chair of Advisory Committee

DEDICATION

To my family.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Rudra Dutta for his guidance and support throughout my doctoral study. I would also like to thank the members of my committee, Dr. George Rouskas, Dr. Mihail Sichitiu, and Dr. David Thunte for their patience and expertise.

To my family, thank you for your love, support, and encouragement. Your support strengthen my resolve and sustained my spirit throughout my journey.

This material is based upon work supported by the GEM Fellowship and the National Science Foundation grants 1111088 and CNS-1402053.

TABLE OF CONTENTS

LIST OF TABLES	vii
LIST OF FIGURES	viii
Chapter 1 INTRODUCTION	1
Chapter 2 Summary Thesis Statement	4
2.1 Related Work	4
2.2 Our Contribution	6
Chapter 3 Enabling Application-Aware Routing within an OpenFlow Network	12
3.1 Research Question	12
3.1.1 Problem Definition and Motivation	12
3.2 Our Methodology	13
3.3 Experiment Design	14
3.4 Experiment Results	14
3.5 Conclusions	15
Chapter 4 Choice-related semantics in a realizable protocol	17
4.1 Research Problem	17
4.2 Our Methodology	17
4.2.1 Message format and Interactions	18
4.2.2 Payment Integration	21
4.2.3 Advertisement Specification and Service Composibility	21
4.2.4 Service Provisioning	23
4.3 Real-world Demonstration	24
4.3.1 Economy Plane Demonstration with Software Switching	24
4.3.2 OPCINet ChoiceNet Intenational Demonstration	25
4.4 Conclusions	26
Chapter 5 Experimental Study Design	27
5.1 Modeling Customer and Provider Economic Strategies	27
5.1.1 Customer Strategies	28
5.1.2 Provider Pricing Algorithms	30
5.2 Modeling Provider Resource Allocation Strategy	33
5.2.1 Design of OPCInet-inspired Study Scenarios	34
5.2.2 Pricing Algorithm Considerations for Hybrid Packet-Circuit Operation	40
5.3 Other Study Design Considerations	41
5.3.1 Traffic Arrival Traces	41
5.3.2 Path Alternative Discovery	41
5.3.3 Numerical Results	42
5.3.4 Network Topologies	42

Chapter 6	Impact of pricing algorithms on network performance	47
6.1	Research Question	47
6.2	Experimental Setup	47
6.3	Experiment Results	48
6.4	Conclusions	48
Chapter 7	Impact of customer choice strategies on network performance	52
7.1	Research Question	52
7.2	Experimental Setup	52
7.3	Experimental Results	53
7.4	Conclusions	54
Chapter 8	Impact of waveband size on strategy interaction	58
8.1	Research Question	58
8.2	Experimental Setup	58
8.3	Experiment Results	59
8.3.1	Topology: 19-node European Optical network	59
8.3.2	Topology: 53-node Hibernia network	60
8.4	Conclusions	60
Chapter 9	Impact of different resource provisioning strategies implemented on a multi-provider choice-based network	66
9.1	Research Question	66
9.2	Experimental Setup	67
9.3	Experiment Results	67
9.3.1	Provisioning strategy effect on 53-node topology	67
9.3.2	Provisioning strategy effect on 19-node topology	68
9.4	Conclusions	71
Chapter 10	Impact of network providers offering customers a single predetermined path choice	75
10.1	Experimental Setup	75
10.2	Background: Path Provisioning Strategies	76
10.3	Results and Discussion	76
10.4	Conclusions	78
Chapter 11	Impact of Exchange Points in a multi-provider choice based network	81
11.1	Experimental Setup	81
11.2	Experimental Results	82
11.2.1	Effect of enabling exchange points	82
11.2.2	Disadvantaged Pricing Algorithm	83
11.2.3	Advantaged Pricing Algorithm	84
11.3	Conclusions	84
Chapter 12	Impact of competition between multiple providers with different topologies	91
12.1	Experimental Setup	91

12.2	Experimental Results	92
12.2.1	Competitive Scenario: Same Pricing Algorithms	92
12.2.2	Competitive Scenario: Different Pricing Algorithms	92
12.3	Conclusions	95
Chapter 13	Conclusions	100
BIBLIOGRAPHY		101

LIST OF TABLES

Table 5.1	General Abbreviations and their full names	28
Table 5.2	Provisioning Table Abbreviations and their full names	28
Table 5.3	Simulation parameters: Customer populations	30
Table 5.4	Simulation Parameters	43
Table 5.5	Network Topologies	43
Table 5.6	European Optical Network coordinates	44
Table 5.7	24-node NSF Network coordinates	44

LIST OF FIGURES

Figure 3.1	Logical Topology	14
Figure 3.2	Physical Topology	14
Figure 3.3	MOS with Traffic Controller 1 activated versus baseline and expedited	15
Figure 3.4	Simultaneous Video Streams from the same media server to the same clients; Expedited Video (left), Best-effort video (right)	16
Figure 4.1	ChoiceNet Message Syntax	19
Figure 4.2	ChoiceNet Speaker GUI	19
Figure 4.3	ChoiceNet Protocol Interactions	20
Figure 4.4	ChoiceNet Demo Topology	25
Figure 5.1	Delay Utility Curves	29
Figure 5.2	Circuit Request Flow	35
Figure 5.3	Resource Adjustment Algorithm	36
Figure 5.4	Free Placement Packet Waveband Sacrifice Strategy	39
Figure 5.5	Single Boundary Packet Waveband Sacrifice Strategy	40
Figure 5.6	European Optical Network topology (19-node European Optical Network with 39 edges)	45
Figure 5.7	NSF Network topology (24-node NSF Network with 43 edges)	45
Figure 5.8	Hibernia topology (53-node Hibernia backbone network with 76 edges)	45
Figure 5.9	24-node Generated network (24 nodes with 39 edges)	45
Figure 5.10	24-node network: NSF Network (blue edges) and Generated network (red edges) with mutual edges colored green	46
Figure 6.1	Running Average Link Utilization	49
Figure 6.2	Circuit Acceptance Rate	49
Figure 6.3	Average Running Average Delay: Budget-limited Cheapest Available policy	50
Figure 6.4	Average Running Average Delay: Budget-limited Maximum Utility (with Utility Threshold) policy	50
Figure 6.5	Total Revenue Generated	51
Figure 7.1	Circuit Acceptance Rate: Five service providers, each with an identical 24NSF topology	54
Figure 7.2	Running Average Link Utilization: Five service providers, each with an identi- cal 24NSF topology	55
Figure 7.3	Average Average Delay: Accepted Customer Population: Five service providers, each with an identical 24NSF topology	55
Figure 7.4	Delay Experienced by Budget-limited Maximum Utility customers grouped by their QoS profile: Five service providers, each with an identical 24NSF topology	56
Figure 7.5	Total Revenue Generated: Five service providers, each with an identical 24NSF topology	57

Figure 8.1	Running Average Link Utilization: Five service providers, each with an identical 19EON topology	61
Figure 8.2	Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology	62
Figure 8.3	Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology	62
Figure 8.4	Total Revenue Generated: Five service providers, each with an identical 19EON topology	63
Figure 8.5	Running Average Link Utilization: Five service providers, each with an identical 53Hibernia topology	63
Figure 8.6	Circuit Acceptance Rate: Five service providers, each with an identical 53Hibernia topology	64
Figure 8.7	Accepted Customers Average Running Average Delay: Five service providers, each with an identical 53Hibernia topology	64
Figure 8.8	Total Revenue Generated: Five service providers, each with an identical 53Hibernia topology	65
Figure 9.1	Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 53Hibernia topology	68
Figure 9.2	Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 19EON topology	69
Figure 9.3	Running Average Link Utilization: Five service providers, each with an identical 53Hibernia topology using Single Boundary Provisioning (compare with Figure 8.5)	69
Figure 9.4	Total PWSePA offered: Five service providers, each with an identical 53Hibernia topology with waveband size of 1λ	70
Figure 9.5	Total PWSePA offered: Five service providers, each with an identical 53Hibernia topology with waveband size of 4λ	70
Figure 9.6	Running Average Link Utilization: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning (compare with Figure 8.1)	71
Figure 9.7	Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning	72
Figure 9.8	Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning	72
Figure 9.9	Total Revenue Generated: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning	73
Figure 9.10	Total PWSePA offered: Five service providers, each with an identical 19EON topology with waveband size of 1λ	73
Figure 9.11	Total PWSePA offered: Five service providers, each with an identical 19EON topology with waveband size of 4λ	74
Figure 10.1	Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 19EON topology	77

Figure 10.2	Running Average Link Utilization: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figure 8.1b)	78
Figure 10.3	Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.3b)	79
Figure 10.4	Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.2b)	79
Figure 10.5	Total Revenue Generated: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.4b)	80
Figure 11.1	Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 24NSF topology and a different pricing algorithm	85
Figure 11.2	Successful Circuit Demands: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points	85
Figure 11.3	Successful Circuit Demands: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points	86
Figure 11.4	Running Average Link Utilization: Five service providers, each with an identical 24NSF topology with an different path provisioning strategy while offering path alternatives using exchange points	86
Figure 11.5	Running Average Link Utilization: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points	87
Figure 11.6	Total Revenue: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points	87
Figure 11.7	Total Revenue: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points	88
Figure 11.8	Experienced Delay: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points	88
Figure 11.9	Experienced Delay: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points	89
Figure 11.10	PWSePA Successfully Offered: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points	89
Figure 11.11	PWSePA Successfully Offered: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points	90
Figure 12.1	Successful Circuit Establishment: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies	93

Figure 12.2	Experienced Delay: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies	93
Figure 12.3	Running Average Link Utilization: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies	94
Figure 12.4	Portion of Selected Packet Waveband Sacrificed enabled Path Alternative: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies	94
Figure 12.5	Total Revenue: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies	95
Figure 12.6	Successful Circuit Establishment: Two service providers, each with a different Pricing Algorithm and different topology	96
Figure 12.7	Successful Circuit Establishment: Two service providers, each with a different Pricing Algorithm and different topology	96
Figure 12.8	Total Revenue: Two service providers, each with a different Pricing Algorithm and different topology	97
Figure 12.9	Total Revenue: Two service providers, each with a different Pricing Algorithm and different topology	97
Figure 12.10	Experienced Delay on Network A when competing against another service provider, each with a different Pricing Algorithm and different topology	98
Figure 12.11	Experienced Delay on Network B when competing against another service provider, each with a different Pricing Algorithm and different topology	98

CHAPTER

1

INTRODUCTION

Computer communication networks have transformed the face of the modern world, by enabling computing to be infused into more and more areas of human endeavor, and computers distributed across the world to engage in cooperative computing tasks in numbers and variety undreamed of just a few years ago. Much of the underlying substrate of the computer networks of the world consist of optical networks. The celebrated layering discipline has allowed the design and implementation of optical communication and networking technology on the one hand, and that of internetworking techniques, protocols, and equipment on the other, to largely proceed independent of each other. However, areas of research that address network global performance characterization and optimization, treating higher layer traffic demands as input, and the realities and capabilities of the underlying optical transport mechanisms as constraints, have been important ones at the junction of these two networking paradigms for the past few decades. In optical networking literature, such studies have often been referred to by the general name of traffic grooming, and come under the general areas of network planning, network engineering, and traffic engineering in the more general context of networking. Traditional mechanisms such as traffic contracts, SLAs, shaping, and policing would be used to limit the entry of traffic into the network, and label-based routing, load balancing, and state-space optimization would be used to enact the traffic and network engineering. The result would be improvement of network performance metrics such as utilization.

More recently, a different thread of network research, in the general area of network economics, has addressed the possibility of fine-grain agile networking services pricing, purchase, and deploy-

ment, as a conceptual parallel to the “App Store” model of application marketing, purchase, and deployment that emerged in the late 2000’s. Drawing further from that model, such research has hypothesized that if customers seeking network services can communicate their requirements at a fine-grain both in the magnitude, and time (say, per circuit request), and receive multiple alternatives from a single or multiple providers, their exercise of choice (each informed by their own knowledge of their utility functions) can create a competitive marketplace in which consistently satisfactory service providers gain reputation and revenue. Further, in such a choice-based networking service model, small providers have an ecosystem in which the services they provide can gain visibility to end-consumers, and be appropriately rewarded, instead of becoming ancillaries to a few large vendors or providers where their contribution is bundled with coarse-grain service contracts. The ChoiceNet project, and other research following on it, have articulated the architectural characteristics of such choice-based networks.

In this thesis, we explore a thread of research on “choice grooming”. This refers to the following idea. Consider a choice-based network model, in which network service providers provide multiple service alternatives for customers to choose from, but also provide estimated performance metrics of interest to the customer accompanying each alternative (say, multiple different path alternatives for the same source-to-destination optical circuit request, with accompanying estimate of path delay), as well as a price quote for each. The price, in some sense, reflects the value that the provider places on the resources that would be utilized to realize the corresponding service alternative, and can reflect not just the opex and capex of the network equipment, but also performance targets. The estimation of the performance impact of establishing each alternative can be dynamic, and reflect the current state of the network, in terms of traffic occupancy, or pattern of residual capacity.

When a customer makes a choice between such alternatives (including the always-available option to choose none; that is, defect), they in turn indicate their perceived utility for the indicated performance metric values by the choice they make. This means that by deciding which service alternatives to offer (in response to a given service request) and how much to price each, the network provider can affect the choices made by the customer, thereby grooming traffic without the traditional traffic engineering mechanisms. Conversely, by providing multiple alternatives, and providing objective performance criteria for each, the network provider enables a diverse population of customers to be potentially more satisfied with (that is, receive more utility from) their received service than if the network had preemptively provided just a single alternative in response to each customer request (traditional grooming). Thus, diversity in customer population, coupled with choice mechanisms, enables both higher overall satisfaction of the customer population, and better network global performance optimization. We refer to this mechanism as *co-optimization* of the network by providers and customers.

The bulk of our thesis consists of studies on the relative impact of various pricing and path selection strategies on the part of providers, and choice preferences on the part of customers. In

particular, we study scenarios inspired by OPCINet, an optical networking technology that enables hybrid optical circuit and packet switching in the same switch, and agile re-apportioning of wavebands for circuit-switching and packet-switching. Before that, we undertake two more pragmatic units of work, to demonstrate that fine-grain manipulation of traffic can be feasibly attained with commercial equipment, and marketplace-mediated interaction between providers and customers can be integrated with OPCINet in a real-world network scenario.

For a combination of non-academic reasons, this thesis work has taken longer than typical to complete. The two preliminary units of work mentioned above now appear dated, but are included for the sake of completeness. We published peer-reviewed papers on both those units, and they have been built upon and improved by others. The core of our work, much of which is more recent, continues to be part of active network engineering research. We have published some of it, and hope to publish the most recent units shortly.

In the following chapter, we present reviews of relevant prior work, and a summary of our contributions in context, preparatory to presenting our contributions in more detail.

CHAPTER

2

SUMMARY THESIS STATEMENT

2.1 Related Work

Traditionally a provider applies some internal strategy in selecting a path for a given customer's connection request (call). Neither the strategy applied nor the cost (to the provider) for the selection are typically exposed to a customer. The customer only knows that their connection request has been carried, in exchange of a price. Depending on the diversity of the customer population and the strategy applied by provider, the network can result with an inefficient use of resource. Our hypothesis is that by using choice as a mechanism we can achieve collaborative optimization of the network for both the provider and customer. [BD14b] provided an early indication that through using such an approach, optimization is possible.

With the **Choice-based approach**, the provider calculates multiple path alternatives for a given customer's connection request and presents the customer with a collection of alternatives. The collection of alternatives only reveals the price for each alternative and does not expose the actual path that these alternatives represent. The price for each alternative reflects some characteristics of the path, based on the current state of the network and the potential projected state of the network after the call is accepted. It is up to the customer to select an alternative based on their own strategy or reject the providers' offering.

The **ChoiceNet project** describes the introduction of architectural entities into the Internet to enable fine-grain economic interactions [RBCDGNW], [BD14c], [WGCDRBN14], [WGCDRBN12],

[BD13], and [BD14a]. These interactions occur within ChoiceNet's 'Economy Plane', a plane that offers the presentation of competing services from various providers, the formation of contracts with each of these providers to satisfy a customer's service requirements, and the verification of performance between contracted services. Transactions occurring within ChoiceNet's 'Economy Plane' would parallel real-world interactions that take place between service providers and customers. Contractual agreement are formed between service providers and customers using a defined set ChoiceNet Interactions after some payment (or suitable consideration are exchanged) has been fulfilled. A successful payment result with the customer receiving a *token* of some kind. The token would be be used authorize access to paid services within the 'Use Plane'. Services may consist of transport services such as path or pathlet service, or some in-flight service.

Within a ChoiceNet environment, programmable switch using Software Defined Networking (mechanisms) such as the **OpenFlow** protocol would enable a ChoiceNet service provider to offer dynamic supply of services based on demand. The OpenFlow protocol provides an open interface that defines the structure and mechanics for the interface between a switch's data plane and control plane. The data plane is responsible with forwarding/dropping traffic as specified by the controller. Through this architecture, network administrators and researchers have a programmable interface to control the behavior of network for Layer 2 - 4 (forwarding and routing) traffic [MABP-PRST08][SNK12]. Using the proposed extension in [GK14], a network operator could include more custom OpenFlow actions, such as introducing IPsec in-flight service which encrypt and decrypt specific traffic within a software switch. A growing interest among vendors in providing programmable data plane is evident in [Int] and [Fin12]. These data plane offer the ability to enable service-aware routing by including deep packet inspection within the datapath. This functionality enables machine learning applications to use network intelligence on traffic patterns to improve network's security and management of the network resources [SPYFGT13]. Each of these approaches could provide a platform for building custom innovative services that could be advertised within ChoiceNet Marketplace.

With no standard north-bound protocol for configuring OpenFlow switch, multiple implementations have been developed by researchers to allow such functionality. A potential solution could be implemented using a modified version of OpenFlow Configuration and Management Proposal (OF-CONFIG). OF-CONFIG allows for the remote configuration and management of multiple OpenFlow data planes. OF-CONFIG specification defines that a OF-CONFIG compliant switch allows for the configuration of an OpenFlow's switch ports, queues, and controller management address [Ope13]. The OpenFlow Logical Switches abstraction resembles FlowVisor[SGYACMP], a network virtualization model, which provided network operators with the ability to slice resources on an OpenFlow switch between different network operators. An extended version of OF-CONFIG protocol could be used to implement how ChoiceNet service provider provision a customer resources.

Path Computation Element (PCE) describes an entity capable of composing a network path from

a network graph based on a requested path's requirement [FVA06]. The path's requirement describes the requested end points (source and destination location) along with the required bandwidth and network constraints such as Quality of Service (QoS) parameters. Service function chaining would also be possible through a PCE [QN15], by offering paths through an ordered list of service functions to perform a specific treatment on the received packets.

There is growing interest to move away from the rigid constraints of fixed-grid Wavelength Division Multiplexing (WDM) optical networks towards more flexible optical networks, which can offer improved resource allocation strategies for today's dynamic traffic demands [TASPCP14]. The shift toward flexible optical networks covers a full range for optical capabilities for today's switches, from improving the transceivers, to optimizing routing assignment techniques. While much attention has been focused on the most extreme generalization, namely an elastic-grid wavelength multiplexing, other generalizations of the optical circuit abstraction are also of practical interest. Optical Packet Circuit Interconnect networks, or **OPCInet** [Har13], is such an example of non-traditional optical networks. With OPCInet, a wavelength switch can move the boundary between wavelength resources for either optical packet switching (OPS) and optical circuit switching (OCS) to be dynamically shifted based on the network's need, enriching the network operator's capability for handling diverse set of traffic patterns. A converged system that applies both the Choice-based network architectural approach with a flexible optical network would allow a provider to apply traffic grooming algorithms based on their current network resource and advertise the most optimal set of service offerings for a customer to choose from [DRB13], further improving performance of the optical network. Such a marketplace would

OPCInet represents a leading example of emerging technology at low layers of optical communication that can nevertheless have profound effects on the capabilities of the network system at a higher layer, and poses both challenges and opportunities for both network service providers and consumers. For posing our research questions regarding customer-provider strategy interactions realized through economy, and their impact on network performance metrics, we take OPCInet as our central basis. The issue of traffic grooming, as well as choice grooming, require fresh study with each such disruptive new communication technology that changes design and operational constraints for network operations.

2.2 Our Contribution

These are the research questions we address in the coming chapters.

Chapter 3: Enabling Application-Aware Routing within an OpenFlow Network

In this unit, we had asked the question of whether we can design a system where traffic can be differentially routed on the basis of information in the application payload within an integrated OpenFlow environment, without significantly degrading the average quality of video traffic traveling through this environment.

This is significant in the context of network co-optimization because traffic engineering policies are limited to only L2-L4 traffic headers with OpenFlow. By enabling routing on the basis of information found in the Application payload/headers (application-aware routing), network operators would gain greater management control over their network to mitigate congestion, while improving the reliability and resiliency of their network. Application-aware routing can improve quality of service (QoS) for many applications by mitigating the effect of packet loss, delay, and jitter (variance of delay)[Int] [Ari13][SMJI13].

We proved the usefulness of this approach by redirecting packets which require deep examination from the OpenFlow switch to a helper box for processing and routing those packets toward their designated path on the basis of the information from the packet's Application header/payload. This was performed on a working prototype we made using commodity hardware.

In succeeding years, with the rapid progress in the field of software defined networking, such a conclusion has become more obvious; [BD20] provides a recent example of benchmarking the performance of software defined networking in SmartNICs. We nevertheless present our contribution, which was timely at the time it was performed, in its appropriate place in this thesis, as Chapter 3.

Chapter 4: Choice-related semantics in a realizable protocol

In this unit, we asked the question what entities and interactions are necessary to realize an economy plane protocol for searching, advertising, and purchasing network services within a choice-based network. We also address the small set of interactions necessary to realize such a protocol.

The choice aspect of a choice-based network refers to a customer being offered multiple choices of services for a given request by querying the economy plane's marketplace. Through the economy plane, customers and service providers have a common meeting ground to advertise, search, and purchase services. Only through a common protocol can multiple entities communicate their objective to each other. By disassembling the necessary entities, we have defined the minimal set of interactions required to enable a choice-based economy plane.

The entities we defined are composed along intuitive roles expected in a network service marketplace ecosystem: Marketplace, Service Provider, and Customer. With these three entities, we have defined a simple protocol with a few interactions necessary to reflect real-life market economics of networking services. These interactions are Discovery, Token request, Listing, Marketplace query, Planner request, Use Attempt, and Negative Acknowledgement. By having fewer interactions to

standardize and inter-operate, it is likely to operate efficiently and be more maintainable for vendors interested in participating in the economy place. We have built a prototype with our set of interactions and demonstrated it at the 23rd GENI Engineering Conference in Urbana-Champaign and at the 2016 IEEE Infocom in San Francisco.

As with the previous unit, this unit of work, timely and valuable when we undertook it, has since been superseded by further developments in semantic networking. Once again, we present it in its place for the sake of completeness, as Chapter 4.

Studying Customer-Provider Strategy Interplay

The subsequent units of work focus on investigating the consequences of the interplay between strategies of providers and customer populations, on network performance, provider revenue, as well as customer Quality of Experience. These form the core value of our thesis, and remain timely at this writing.

Although different strategies, and study scenarios, are logically encountered at various different units, many are referred to in multiple units of work. For ease of reference, we have gathered them together in one preliminary chapter on the design of our overall study, Chapter 5, before embarking on the chapters presenting the actual studies as follows.

Chapter 6: Impact of pricing algorithms on network performance

In this unit, we ask the question: All other things remaining the same, how does the pricing algorithms (that are used to assign prices to specific offered path alternatives by a provider) impact the performance of the choice-based network? This forms a baseline for the rest of our work.

This question is motivated by the desire to perform network engineering on the part of the provider. A choice-based network provider that employs pricing algorithms may be able to groom demand and maximize their network utilization, if pricing algorithms have an effect on network performance. Grooming demand involves offering some network service choice as more attractive than others; the network provider can encourage users to buy services that are in surplus instead of those which use resources that are viewed as more precious due to their scarcity. This requires that network providers use reasonable strategies in pricing their network resources.

We propose one baseline pricing strategy, in which the price of an offered alternative is strictly proportional to the network resources that would need to be dedicated to realize this alternative (a typical opex + amortized capex model). We also propose two pricing approaches that attempt to reflect the current occupancy of network resources, and consequent relative desirability of realizing each offered alternative from the point of view of the network provider.

We also postulate two different realistic customer strategies for deciding which one of the offered alternatives to choose. We assume newly arriving customers may equally likely use one or the other

strategy, and find that their relative experiences under the above pricing strategies are different.

Chapter 7: Impact of customer choice strategies on network performance

In this unit, we investigate the question on whether the nature of the customer population (with respect to strategies for making choices between offered alternatives) have an effect on the network performance, in turn. This is important because if customer population mix has no discernible effect on network performance, then the interaction between customers and network providers can ignore such population differences in attempting to maximize network utilization.

We have studied the effect of three different compositions of the two previously defined customer choice strategies with various well-accepted network topologies from literature in terms of generated revenue, the network utilization, and the number of customers served. Results from our simulations show that with all other things remaining the same, a different in choice strategy can result in different network performance.

Chapter 8: Impact of waveband size on strategy interaction

In this unit, we repeat the study of the previous chapter with differing waveband sizes, to investigate how the granularity of packet/circuit provisioning affects the strategy interplay. Previously, we had focused on a waveband size of four wavelengths. In this chapter, we consider a “waveband size” of one wavelength; that is, every wavelength can be independently allocated to packet or circuit traffic, and independently switched.

We find, as expected, that the granularity does affect the absolute performance outcomes; but also found that in some cases it appears to change the relative performance of different strategies.

Chapter 9: Impact of different resource provisioning strategies implemented on a multi-provider choice-based network

In this unit, we consider the question of the effect different provisioning strategies have on a multi-provider choice-based network, all other things remaining the same, in terms of the network’s acceptance rate and network utilization. From the customer’s perspective, we also consider whether the customer population also receives some relative benefit of the provider’s provisioning strategy.

A network provider aware of the performance of their resource provisioning strategy with their implemented pricing algorithm can select the ideal policy for a given customer population, network topology, and/or expected competitor’s pricing algorithm. Resource provisioning strategies can dynamically adjust a network link’s waveband assignment between optical circuit switching (OCS) and optical packet switching (OPS). We explore two resource provisioning strategies, Free Placement and Single Boundary, and compare their effect on competing network providers.

Results from our simulation studies show that the resource provisioning strategy had an effect on the performance of multiple provider's network. Depending on the network topology, the competitive ecosystem of the network providers resulted in one resource provisioning strategy out performing another.

Chapter 10: Impact of network providers offering customers alternative choices in a multi-provider choice-based network

In this unit, we examine whether each provide offering customers multiple path alternatives positively or negatively affects the network performance when compared to a scenario where each provider peremptorily offers only a single path alternative, selected by some criteria pre-determined by each provider, in response to each customer request (so that the competition is entirely between different providers, and not between multiple alternatives offered by the same provider).

This is important because providers presenting only heuristically selected paths may result in under utilizing network resources. Several customers may reject the provider's pre-selected path due to cost or delay, despite the provider having a path alternative that the customer would have selected had it been advertised. The interaction between the heuristic and pricing algorithm may result in only long paths being the preferred selection, which may lead to a higher rejection rate from delay sensitive customers.

Chapter 11: Impact of Exchange Points in a multi-provider choice-based network

In this unit, we evaluate the effect of allowing a limited set of network switches to serve as exchange points between the networks of multiple network providers within a choice based environment. Exchange points enables network providers to collaborate with each other to provide customers with a path service.

Cooperating network providers can reduce the total amount of resources occupied to provide a service and thus increase their available resources for future customers, potentially allowing the network providers to serve more customers and generate more revenue. From the point of view of the customer population, cooperating network providers are seen to be able to provide a larger selection of path choices within the marketplace at different price points and different delays. By measuring the networks' performances, we can evaluate the effect of services enabled by cooperation via exchange points.

Chapter 12: Impact of competition between multiple providers with different topologies

In this unit, we evaluate a multi-provider scenario consisting of two competing service providers offering network paths for a shared pool of customers. Each topology contain the same set of

co-located network nodes but consist of unique set of fiber connections between their network nodes.

Pricing algorithms groom the network's performance by making abundant network resources more attractive than paths that rely on scarce resources. This is important, because pricing algorithms allow a service provider to remain competitive even when their available path alternatives are less attractive because they contain higher delays (i.e. longer paths). This allows service providers to balance out their network's congestion and reduce the chance of blocking future customers.

CHAPTER

3

ENABLING APPLICATION-AWARE ROUTING WITHIN AN OPENFLOW NETWORK

3.1 Research Question

Can a system be designed where traffic is differentially routed on the basis of characteristics only found within the application's payload, without significantly degrading the average quality of traffic traveling through this environment? This is the problem we will look at in this chapter. To investigate this, we consider OpenFlow, a networking protocol that enables rapid development for monitoring and routing traffic using a programmable interface, and measure the result of our study by specifically analyzing the quality of transmitted video over the system because video quality is directly impacted by network latency.

3.1.1 Problem Definition and Motivation

OpenFlow is the leading Software Defined Network (SDN) specification and provides an open programmable interface between the network's control plane (the intelligence) and the data (forwarding) plane and allows for the integrated management of Layer 2 - 4 traffic [MABPPRST08][SNK12]. This

allows the control plane to be separated into a different network device than the forwarding engine, allowing network operators to have the flexibility to program the behavior of their network deployment.

Due to the wide diversity in payload structure, the designers of OpenFlow did not include the ability to issue L7-based policies. The obvious approach to expand OpenFlow into allowing application-aware routing (L7-based traffic engineering), would be to enhance the datapath flowtable, or to utilize the controller for deep packet inspection, but both of these options introduce significant scalability barriers. We proposed and prototyped an enhancement to OpenFlow based on the idea of an External Processing Box (EPB) optionally attached to forwarding engines; however, we use existing protocol extension constructs to control the EPB as an integrated part of the OpenFlow datapath. This provides network operators with the ability to make Application specific policies to control service insertion and traffic steering, without breaking the open paradigm. This novel, yet eminently practical augmentation of OpenFlow provides added value, making it critical for realistic networking practice. Retention of multi-vendor openness for such an approach has not been previously reported in literature to the best of our knowledge. In [UD14], we reported the numerical results from our prototype, characterizing the performance and practicality of this prototype by implementing a video reconditioning application on this platform.

3.2 Our Methodology

Our architectural approach embraced both the SDN and Network Function Virtualization (NFV) paradigm to achieve application-aware routing. NFV refers to paradigm decoupling network middle-box operations such as deep packet inspection, intrusion detection system, load balancers, and etc. from hardware to software. This decoupling would allow for flexible provisioning of services while reducing equipment and power cost [Chi12]. Our architecture uses SDN to control and manage the deployment of traffic steering through various virtual network functions. We realize this architecture through an External Processing Box (EPB), a commodity box that performs additional processing on packets sent to it by a datapath.

Using constructs provided by OpenFlow, we extended the OpenFlow Experimenter action to enable application-aware routing through an attached EPB. This OpenFlow action instructed the OpenFlow switch to forward matched L2-L4 traffic to an specified EPB. Any flow that matched that action's match, caused the OpenFlow switch to signal to the EPB on what payload inspection policy to apply on that flow. Specifically, we instructed the EPB to perform video reconditioning based on the video flows' streaming URL. Successful URL matches are routed out with a VLAN tag designating expedited treatment otherwise best-effort. The EPB used the outermost VLAN tag to inform the datapath how to route the traffic. The VLAN tag value is mapped to outgoing port in our experiment. The datapath pop's the VLAN tag before forwarding the traffic out the designated port.

3.3 Experiment Design

Our prototype simulates an environment where requests and responses between a client and media server from streaming video must travel through our OpenFlow network. Figure 3.1 illustrates a logical topology map for this environment. Traffic propagating through the best-effort routes was treated with variable amount of delay. One of these clouds is a Traffic Controller, as seen in Figure 3.2. The Traffic Controller introduces outgoing traffic delay with a variance of $100\text{ms} \pm 50\text{ms}$. Delays on networks are rarely uniformly distributed, so to simulate real-life network delay we have programmed our Traffic Controllers to have a normal distribution of variation in delays [Net].

Figure 3.1 Logical Topology

Figure 3.2 Physical Topology

3.4 Experiment Results

We used Video-Tester [UNOALS12], a video testing platform, to objectively evaluate video quality using Peak Signal Noise Ratio based Mean Opinion Score (PSNR-MOS) and Structural Similarity Index (SSIM) to measure the perceived video quality. Videos used for evaluating this architecture are among the videos sequences used predominately in video research [Res]. Figure 3.3 illustrates the video quality assessment of the experiment when Traffic Controller 1 is activated. Best-effort traffic maintained a similar PSNR-MOS regardless of whether the traffic was routed through the EPB or not. The lack of significant difference between these values illustrates that the inclusion of the EPB into the circuit, even when it is not expediting traffic, had little effect on the video quality. As expected PSNR-MOS was at its highest with the baseline, which the expedited traffic closely mirrored. A screen shot of two simultaneous video feeds being actively reconditioned by our setup is illustrated in Figure 3.4. Note that the video feeds originated from the same media server (see application's header) and was received by the same client, yet the video on the left is being received without any visually distortion. In this case, the video received expedited treatment because the

video's URL contained "channel7".

Figure 3.3 MOS with Traffic Controller 1 activated versus baseline and expedited

3.5 Conclusions

We evaluated the EPB architecture within a OpenFlow network and measured the QoS metrics and captured video quality assessment measurements for video traffic streaming through it. With the EPB optionally attached to a forwarding engine, network operators can engage in rapid development cycles, reduce capital and operating expenditures (CapEx/OpEx)[Hoe12], and create routing policies based on the application's anticipated bandwidth requirement, malicious payload detection, or content-availability [Fin12]. Most important, based on the results of our study we have shown that EPB can be attached to a forwarding engine without sacrificing performance. We have published

Figure 3.4 Simultaneous Video Streams from the same media server to the same clients; Expedited Video (left), Best-effort video (right)

our research contribution presented in this chapter in [UD14].

CHAPTER

4

CHOICE-RELATED SEMANTICS IN A REALIZABLE PROTOCOL

In this chapter, we describe a semantic language we developed for enabling an economy plane within ChoiceNet, a choice-based network. Through interactions within the economy plane, customers are able to search, advertise, negotiate, and purchase network services.

4.1 Research Problem

To compose an economy plane for a choice-based network, a set of entities and interactions are necessary to realize a protocol capable of enabling users to search, advertise, and purchase network services within a choice-based network. This chapter will address the system research question of identifying and realizing a small set of interactions sufficient to realize such a protocol, and demonstrating its sufficiency by practical implementation.

4.2 Our Methodology

We have defined three distinct entities: Marketplace, Service Provider, and Customer. With these three entities, we have defined a simple protocol with a few interactions necessary to make contractual agreements between service providers and customers.

4.2.1 Message format and Interactions

We have defined a set of messages necessary to complete a minimum set of interactions between a customer and service provider for contracting the use of services. Each interaction defines a protocol for handling its objective. Each ChoiceNet message consists of a set of fields, individually called ChoiceNet Message Field, that contains 3-tuple structure; *Attribute Name*, *Value*, and *Vocabulary URL*. The *Attribute Name* describes the field's header and the *Vocabulary URL* describes the location for the field's definition. The *Value* can contain a literal or another ChoiceNet Message Field(s), allowing it to be hierarchical. A ChoiceNet Message contains the following ChoiceNet Message Fields Version, Originator Name, Originator Signature, Originator Type, Message Type, and Message Specific Field. Figure 4.1 illustrates a basic ChoiceNet Message. Version defines the ChoiceNet Semantic Language version being used by the message. Originator Name, Originator Signature, Originator Type describe the entity's name, security signature, and the ChoiceNet Provider type, respectively. Message Type describes its type and Message Specific has the payload for that message type. Each ChoiceNet messages are encapsulated within an Extensible Markup Language (XML) payload and are currently transported within an User Datagram Packet (UDP), though any transport protocol would be permissible. Figure 4.2, shows a configurable Java applet 'ChoiceNet Speaker' that we have developed to allow a user to participate as any ChoiceNet-entity (service provider or customer) and perform any of the defined interactions.

- Rendezvous Interaction
- Purchase Service Interaction
- Advertise Service Interaction
- Marketplace Interaction
- Planner Interaction
- Negative Acknowledgment Interaction
- Use Attempt Interaction

Figure 4.3 illustrates the handled transactions, with the mentioned interactions. The mapping for the transactions in Figure 4.3 can be found in the following listing. In this example, it is assumed that each entity initializing communication already knows the recipient entity and what role that entity is serving through a Rendezvous message. This Rendezvous message is not shown in the figure.

1. Token Request: **Purchase Service Interaction** for a Listing Service

- Token Response/NACK: Positive/Negative Acknowledgment regarding the listing service during their **Purchase Service Interaction**

Figure 4.1 ChoiceNet Message Syntax

Figure 4.2 ChoiceNet Speaker GUI

Figure 4.3 ChoiceNet Protocol Interactions

3. Listing Request: **Advertise Service Interaction**'s listing an advertisement
4. Listing Confirmation/NACK: Positive/Negative Acknowledgment regarding **Advertise Service Interaction**'s attempt in listing an advertisement
5. Marketplace Query: **Marketplace Interaction**'s for a Marketplace search
6. Marketplace Response: Results from **Marketplace Interaction**'s Marketplace Search
7. Token Request: **Purchase Service Interaction** for some service
8. Token Response/NACK: Positive/Negative Acknowledgment regarding the service during their **Purchase Service Interaction**
9. Use Attempt: The **Use Attempt Interaction** for activating and provisioning a service
10. Use Attempt Confirmation/NACK: Positive/Negative Acknowledgment regarding the **Use Attempt Interaction**

The overhead of building a service contract between customer and service provider requires at least six messages to be successfully transmitted and retrieved within the 'Economy Plane' to search the Marketplace, purchase, and use a service. Services that last years, months, or a few minutes can be practically governed as a choice service, but not services with lifespans of seconds or a few packets. The overhead of building a service contract should not be the majority of the overall traffic. Assuming each ChoiceNet message fits in single packet, a 5% overhead can be maintained by any service that engages in at least 120 total packets.

Due to the fluidness within a choice-based network, all entities can behave as either the customer or provider. This can be demonstrated in an eco-system with at minimum three entities: a

marketplace, customer, and a single third-party service provider, as illustrated in Figure 4.3. The service provider offers some service, such as a pathlet or transformation service. The Marketplace maintains a repository of service advertisements for which it and other service providers can list their service offerings. In order for service providers to advertise a service within a Marketplace, the service provider must retrieve a token from the Marketplace before they are allowed to advertise their service. In this case, the service provider is a customer of the Marketplace. When a customer discovers an advertised service that they wish to purchase, the customer contacts the service advertisement's owner to begin the Purchase Service Interaction.

4.2.2 Payment Integration

Customers must pay for an advertised service before they are allowed to utilize its service. To enable this, we have integrated our ChoiceNet implementation with payment transactions system using third-party electronic payment service's (PayPal and Coinbase's Bitcoin Testnet) sandbox environment. Each ChoiceNet entity maintains a payment portal, which can be used to interact directly with another ChoiceNet entity's payment portal. The third-party electronic payment service provides a Transaction ID with a successful payment, which is later used during the Purchase Service Interaction to request a token for service use. The requested token is required to be used for any service provisioning (Advertise Service Interaction or Use Attempt Interaction), as it signifies a proof of purchase or proof of an outstanding commitment of service.

4.2.3 Advertisement Specification and Service Composability

To facilitate marketplace alternatives and enable fine-grain service composition, a service specification must represent a general-extensible design for describing a service. By default, an advertisement specification includes the service name, service detail, description, provisioning parameters, the provider's name, price, and provider's purchase portal information. Advertisement specification are represented within an XML file, to allow easy adoption and extensibility [UD15]. Services such as measurement, analytic, in-path service, verification, monitoring, are possible to represent within this semantics. Listing 4.1 illustrates an example of a transit service specification. When a service specification is entered into a marketplace it is viewed as a service advertisement by other ChoiceNet entities.

The only interaction not mentioned in Figure 4.3, is the Planner Interaction. The purpose of this interaction is to allow a customer, who is unable to find an advertised service that matches their service requirement, the ability to have a composition (planner) service provider compose service recipe from multiple service providers that does resolve the customer's service requirement. This interaction is design to use the `PLANNER_REQUEST`, `PLANNER_RESPONSE`, and `NACK` messages. `PLANNER_REQUEST` informs a planner provider of the service requirement a customer requests.

```

<advertisement>
  <service>
    <name>Transit Service to Wayne Enterprise network</name>
    <details>
      <source_location scheme="IPv4" value="10.1.10.1/32" />
      <destination_location scheme="IPv4" value="10.1.12.0/24" />
    </details>
    <description>Transit Service Advertisement</description>
  </service>
  <price method="USD" value="21" />
  <providerID>ACME Corporation</providerID>
  <purchasePortal>
    <source_location scheme="Payment" value="Bitcoin:
      mknFpFW8x5pvLH8WSSLSKQRBrkPPiPPoxFv" />
  </purchasePortal>
  <provisioningParameters>
    <configurableParameters/>
  </provisioningParameters>
</advertisement>

```

Listing 4.1 Transit Specification

PLANNER_RESPONSE returns a list of feasible service compositions that satisfies the customer's request, while NACK informs the customer/provider of any error that may have occurred during their interaction.

The planner uses the service advertisement as a service primitive to compose feasible service recipes from third-party service providers' advertised services. Advertised services stored within a Marketplace, provides a planner service with the available network resource and functionality being offered from a host of third-party service providers. These advertised services must adhere to some standardized semantics to allow service compositions between multiple third-party service providers. We have investigated and prototyped a minimum example on Global Environment for Networking Innovations (GENI) [UD15][Ell08]. A planner's service composition contains the Advertisement ID and the provisioning parameter necessary to connect one service to another. Advertisement ID identifies the advertisement within the marketplace, while the provisioning parameter informs the customer how they should tell the provider to route their traffic. Listing 4.2 provides an example of what a feasible recipe looks like within a PLANNER_RESPONSE message. In this example, the provisioning parameter gives the next hop address a customer should have their traffic forwarded towards within a service provider's network. The provisioning parameter is intended to be a versatile field to pass-along any network forwarding configuration, for example it could be envisioned to assist in passing a lightpath's wavelength assignment between multiple optical network providers for a customer's flow [UD15].

Due to the close relationship between the Planner service and the Marketplace, the two providers were envisioned as a single entity (Marketplace) when designing the ChoiceNet semantic protocol. The Marketplace Interaction is viewed as a free query for a matching advertised service, while the Planner Interaction provides a composed service recipe(s) for a service request.

```
<recipe>
  <cost value="71" />
  <serviceAdvertisements>
    <advertisement identifier="08e19dc233f20ca274dd1c2270003dca"
      provisioningParameter="10.1.10.1/32" />
    <advertisement identifier="08e19dc233f20ca274dd1c22700050df"
      provisioningParameter="10.1.12.2/32" />
    <advertisement identifier="08e19dc233f20ca274dd1c2270006875"
      provisioningParameter="10.1.2.50/32" />
    <advertisement identifier="08e19dc233f20ca274dd1c22700085df"
      provisioningParameter="10.1.23.3/32" />
  </serviceAdvertisements>
</recipe>
```

Listing 4.2 Service Composition Recipe

4.2.4 Service Provisioning

Provisioning is an internal process performed by a service provider and takes place within the provider's Control Plane. In our current implementation service provisioning can take place in two ways, either for advertising a service or for a general service. With both approach, the token received after purchasing a service is required along with either an advertisement or a firewall specification. Advertising a service uses the LISTING_REQUEST message with an advertisement specification, while general service provisioning uses the USE_ATTEMPT message along with a firewall specification. The firewall specification informs the provider what traffic characteristics should be allowed through its network and how it should be routed (next hop). Specifically, the traffic characteristics describe the appearance of the packet headers, while the provisioning parameter describes the next hop location. A service provider provided with a valid token and the necessary specification, will provision the service. For advertising a service this involves uploading an advertisement specification into the Marketplace's advertisement repository, while for a general service it involves activating a service for a desired traffic match. Any attempt to provision service without an appropriate token or service payment, results in a Negative Acknowledgement being received.


```

<firewallSpecification>
  <action>ACCEPT</action>
  <addressVersion>IPv4</addressVersion>
  <destinationAddress>ANY</destinationAddress>
  <destinationPort>ANY</destinationPort>
  <protocol>ANY</protocol>
  <sourceAddress>10.1.10.1/32</sourceAddress>
  <sourcePort>ANY</sourcePort>
  <provisioningParameter>10.1.23.3</provisioningParameter>
  <previousNetwork>10.1.2.50</previousNetwork>
</firewallSpecification>

```

Listing 4.3 Use Plane Firewall Specification

4.3 Real-world Demonstration

4.3.1 Economy Plane Demonstration with Software Switching

We have explored advancing our prototype to a deployment that fully realizes the formation of an ‘Economy Plane’ contract and allows for the propagation of data within that contract’s ‘Use Plane’ resources, after a contract has been finalized. This prototype also allowed the verification of a service composition. The prototype consisted of four service providers, a Marketplace, customer, and a third-party content server. Among the service providers, a single provider is solely provides a composition (planner) service, while the remaining providers demonstrate network providers offering pathlet services. An intermediate network service provider also offers transformation services, specifically a substring-specific payload modification and packet logging service. The topology for this system is illustrated in Figure 4.4. The goal for this deployment was to both validate that by using the current set of ChoiceNet Interactions, a customer could successfully send traffic through the network service providers to a connected content server and verify that our composition service could compile feasible service recipe with the given information in the Marketplace. To demonstrate the efficacy of the prototype, it was deployed on GENI with each of its entity and data plane resources within separate virtual machines [BUDR16].

Each network provider maintained a programmable Software Defined Network (SDN) switch, specifically the open source software switch from FlowForwarding called Link Is Not Closed (LINC). SDN switch allows granular control over the network using an out-of-band third party controller software. When a customer sent a valid request to activate the provisioning of their service, the provider’s economy plane agent would send the firewall specification (along with the valid token) to the SDN controller. The controller assessed the specification and compiled a corresponding OpenFlow rule based on the content of the specification. The time span of the rule is based on the expiration value of the token. The service provider offering a transformation service used a simpler

Figure 4.4 ChoiceNet Demo Topology

version of the External Processing Box (EPB) (discussed previously in Section 3) to perform the substring-specific payload modification and packet logging. Each service provider advertised their suite of services within the Marketplace for customer to discover and potentially purchase. With this deployment and the set of services, we demonstrated all possible end-to-end services with the combinations of advertised services; pure transit services, transit services with payload modification, transit services with packet logging, and a transit service with both payload modification and packet logging.

By default traffic sent to a network service provider’s ‘Use Plane’ network prior to a Purchase Service Interaction or prior to an Use Attempt Interaction occurring within the ‘Economy Plane’, are dropped by the network. Only after these interactions are successful and a firewall specification is supplied that appropriately matches the customer’s traffic headers, will the traffic be allowed to traverse the provider’s network.

4.3.2 OPCINet ChoiceNet Intenational Demonstration

OPCINet, which we described in Chapter 2 and which provides the motivating basis of our network model for the research presented in subsequent chapters, consists of optical switches capable of both Optical Circuit Switching, and Optical Packet Switching, and dynamically allocating wavebands to one or the other use. The decision to switch a waveband to and from circuit/packet utilization would be based on market forces of demand and supply in the choice-based view of network co-optimization. In an effort to validate the practical feasibility of such an integration of our ChoiceNet interactions implementation and a real-world industrial system such as OPCINet, we successfully created a “hero experiment” in which the control and management of the OPCINet network in

NICT's Tokyo Labs was provided by ChoiceNet agents, and coordinated in real-time through a secure connection to our lab at NC State University where those agents as well as the Marketplace implementation was running.

4.4 Conclusions

We have defined a set of entities that are composed along intuitive roles expected to enable a network service marketplace ecosystem. With these entities, we have defined a simple protocol with a few interactions necessary to reflect real-life market economics of networking services. These interactions include the ability to search, advertise, and purchase network services. By having fewer interactions to standardize and inter-operate, it is likely to operate efficiently and be more maintainable for vendors interested in participating in the economy place. We have published our research contribution presented in this chapter in [UD15], [UBDR16], and [UDMH18b]. We documented our open-source proof-of-concept implementation at the 23rd GENI Engineering Conference in Urbana-Champaign and at the 2016 IEEE Infocom in San Francisco in [BUDR16] as mentioned before, and we also reported on our OPCINet implementation in [MHUD16; MHUD18].

Beyond this work, we also presented a potential extension of the same to the context of 5G networks in [BUDR17a], and in [BUDR17b] addressed the issue of dynamic real-time composition of network services. These were primarily the work of fellow doctoral researcher Shireesh Bhat, and we do not report on the details of these units of research in this thesis.

CHAPTER

5

EXPERIMENTAL STUDY DESIGN

As we remarked before, our thesis research focuses largely on studying the effect choice has on the performance optimization of networks that provide a marketplace of network services. In this preliminary chapter, we describe our methodology for, and define the various provider/customer strategies, network topologies, and other study parameters of, our studies presented in subsequent chapters.

5.1 Modeling Customer and Provider Economic Strategies

In this section, we report on a set of strategies that we applied on customers exercising choice and a set of pricing strategies executed by a service provider(s) within choice-based network. The Choice-based network grants its users with the ability to exercise choice from a list of advertised services from the provider(s), unlike a traditional architecture where a provider preemptively allocates resources based on their internal strategy. A user (or a customer) may use different policies when selecting a service. While a provider may use a pricing strategy to value the preciousness of network resources to achieve optimal network performance. Our hypothesis is that the interplay of these strategies will result in better utilization of the network and better overall value for both customer and provider. This concept is what we call co-optimization, optimization through collaboration. Each of these strategies are discussed in detail in the following subsections.

Table 5.1 General Abbreviations and their full names

Abbreviation	Full Name
OPCInet	Optical Packet and Circuit Integrated Network
PPA	Passive Pricing Algorithm
WCRPA	Weighted Critical Resource Pricing Algorithm
NBPPA	Network Blocking Projection Pricing Algorithm

Table 5.2 Provisioning Table Abbreviations and their full names

Abbreviation	Full Name
FP-PWSePAoP	Free Placement with PWSePA offered persistently
PWSePA	PWS enabled Path Alternative
PWSePAoP	PWSePA offered persistently
RAA	Resource Adjustment Algorithm
SB-PWSePAoP	Single Boundary with PWSePA offered persistently

5.1.1 Customer Strategies

5.1.1.1 Budgetary Constraint and Defection

The customer population is composed of budget-limited customers with a path selection preference of selecting either the cheapest advertised path or the path with the largest delay-utility greater than the customer's utility threshold. Budget-limited customers are sensitive to price and only select paths' priced below their budgetary constraint. The budgetary constraint reflects a fractional premium a customer is willing to pay over a reasonable cost estimate for the given source-destination before defecting. We assume that customers are aware of an "expected normal cost" for a circuit between every source-destination pair; this assumption reflects the common situation in competitive markets that prices charged and paid become community knowledge. In our study, we use the cost of the shortest path for this cost estimate. Customers use this knowledge to make decisions to accept some alternative versus defecting; thus if an estimated cost for a path is X and a customer's fractional premium is $.8$, then the customer's threshold would be $X + (X * .8)$, and the customer would defect if and only if none of the offered alternatives had a price below this threshold.

In our studies, we assumed each newly arriving customer is equally likely to have a fractional premium of 0.25, 0.5, 0.75, or 1.

5.1.1.2 Utility-based Choice

Utility functions serve as a mathematical representation of the customer's preferences and provides a mathematical tool for comparing bundles, specifically for customers interested in their lightpath's delay. We based our delay-utility curves (η) on the research discussed in [MMG08] and randomly

applied one of the application profile's delay utility curve to each delay-utility customer. A breakdown of the curves' functions are indicated in Equations 5.1 - 5.3. Each mathematical functions in a utility curve are color differentiated and is illustrated in Figure 5.1. Delay-Utility customers gauge the utility of a lightpath based on its normalized delay ranking among alternative paths with the same source and destination. All delay-utility customer have a randomly assigned minimum delay utility threshold they are willing to accept. A delay-utility customer will then select the lightpath that provides them with the highest utility above their delay-utility threshold (and below their budgetary constraint).

$$\eta_{VoIP} = \begin{cases} 100 - .070756/x & 0 \geq x < 3/56 \\ -25 \tanh(.04 * (x - 205)) + 65 & 3/56 \geq x \leq 3/28 \\ 44.841 - (.016 * x) & 3/28 > x \leq 1 \end{cases} \quad (5.1)$$

$$\eta_{Video} = \begin{cases} 100 - .004 * x & 0 \geq x < 1/14 \\ -25 \tanh((x - 625)/200) + 75 & 1/14 \geq x < 5/14 \\ 57 - (.006 * x) & 5/14 \geq x \leq 1 \end{cases} \quad (5.2)$$

$$\eta_{OnlineGaming} = \begin{cases} -45 \tanh(.015 * (x - 1.35)) + 60 & 0 \geq x \leq 1 \end{cases} \quad (5.3)$$

Figure 5.1 Delay Utility Curves

Table 5.3 Simulation parameters: Customer populations

v	Customer Population Distribution	
	25%	Budget-limited Cheapest Available policy
	75%	Budget-limited Maximum Utility (with Utility Threshold) policy
	50%	Budget-limited Cheapest Available policy
	50%	Budget-limited Maximum Utility (with Utility Threshold) policy
	75%	Budget-limited Cheapest Available policy
	25%	Budget-limited Maximum Utility (with Utility Threshold) policy

Our studies focused on three different customer population ratios of cheapest-priced and delay-utility-sensitive customers. We studied customer populations with three different ratios of these two types of customers, as shown in Table 5.3; each newly arriving customer is either cheapest-priced or delay-utility-sensitive with these relative probabilities. This characteristic of the customer is independent of their random budget constraints (that we described earlier in this section).

Although we have only studied circuit alternatives with different customer utilities based on delays, this serves to illustrate the general approach of the provider providing performance estimates on offered alternatives, and the customer making a choice informed by their utility functions applied to the same, in addition to the price. For example, providers could provide alternatives that would be carried on specific (perhaps old) fiber bundles, and indicate estimated non-linearities or other impairments for each. Customers who depended on highly linear operation, because of analog use or aggressive coding of digital signals, would use a different utility function to guide their choice decisions than customers non-sensitive to such impairments.

5.1.2 Provider Pricing Algorithms

The provider's pricing algorithm provides a framework for computing prices for each network resource (or feasible path). An advertised alternative must be reasonably priced. The higher a price may suggest that the provider wants the customer to use it less, or that customer may want to use it more. In a non-manipulative Choice-based system, these two objectives should align. The following are the studied pricing algorithms:

- Passive Pricing Algorithm (PPA)
- Weighted Critical Resource Pricing Algorithm (WCRPA)
- Network's Blocking Projection Oriented Pricing Algorithm (NBPPA)

5.1.2.1 Passive Pricing Algorithm

Passive Pricing Algorithm is a flat-rate pricing based on the base cost of the path's resources. For example, if the bit rate along a link is priced at X cents per Mbps and the switching capacity is Y cents per Mbps, then a path that transmits over three links with two intermediate switches would be the sum price of the three links, the two intermediate switch, and usage of the switching capacity from both the source and destination node, i.e. $3*(X)+2*(Y)+2*(Y)$. If the link cost were 12 per unit (cents, dollars, pounds, etc.) and switching cost going at 2 per unit, then the total cost of the path would be $12*3+2*4 = 40$ units.

5.1.2.2 Weighted Critical Resource Pricing Algorithm

Weighted Critical Resource Pricing Algorithm (WCRPA) is a responsive strategy that attempts to dynamically shift the cost of resources based on the network's perceived scarcity in resources. The intuition behind this approach is that as popular links become congested and start becoming bottlenecks, path provisioning algorithms will start using alternate links, which will be the next bottlenecks, and form a steadily reducing set of key resources in maintaining network connectivity. Those secondary links should at that point be used more sparingly than before. In the economic framework, the mechanism for encouraging such sparing use is by price adjustment. The algorithm works by translating the network's residual capacity graph into a strongly connected components (SCC) graph using Tarjan's Algorithm [Tar71], where links with residual capacity below a threshold are removed prior to running the Tarjan's Algorithm on the network graph. Links connecting different SCCs (a SCC border resource) are penalized with a premium (surcharge) as these resources are now perceived as scarce due to their connectivity, while resources within a given SCCs are discounted. The cumulative sum of penalties and discount along a path including the path's base cost provides the final cost of the path. Algorithm 1 describes how WCRPA applies pricing penalties on resources deemed critical.

In the above, the discounting step is performed in order to normalize away the normal scarcity pricing approach in common economic interactions. In other words, we model a provider who uses relative pricing of resources entirely as a means to signal relative scarcity of resources to a potential customer (that is, the relative desirability of using different resources from the provider's point of view of most beneficial network operation), and not to perform "price gouging", that is unfairly seize opportunities for increasing profit as supplies diminish and demands do not.

We investigate three different thresholds with this pricing algorithm: 25%, 50%, and 75%.

5.1.2.3 Network Blocking Projection Pricing Algorithm

Network Blocking Projection Pricing Algorithm (NBPPA) is a responsive strategy that prices paths based on the blocking projection a path has on the current network's available circuit resources.

Algorithm 1: Weighted Critical Resource Pricing Algorithm

Input : Network Nodes $\mathcal{N}$, Links $\mathcal{L}$, $\mathcal{G}=(\mathcal{N}, \mathcal{L})$, Current Utilizations $U(l), \forall l \in \mathcal{L}$

```
1 scarcityThreshold = getScarcityThreshold()
2 congestedLinks  $\mathcal{L}_g = \{l \in \mathcal{L}, U(l) > (1 - \textit{scarcityThreshold})\}$ 
3 scc = getTarjanSCC( $\mathcal{G}'(\mathcal{N}, \mathcal{L} - \mathcal{L}_g)$ )
4 # avoid extreme cases of all links getting only rebates/premiums
5 if scc.size > 1  $\wedge$  scc.size  $\neq$  N.size then
6   | criticalLinks  $\mathcal{L}_c = \{l \in \mathcal{L}, \text{head and tail of } l \text{ are in different components of } \textit{scc}\}$ 
7   | noncriticalLinks  $\mathcal{L}_n = \mathcal{L} - \mathcal{L}_c$ 
8   | totalBW =  $\mathcal{L}.\textit{totalBandwidth}$ 
9   | premium =  $\mathcal{L}_c.\textit{availBandwidth} / \textit{totalBW}$ 
10  | rebate =  $\mathcal{L}_n.\textit{availBandwidth} / \textit{totalBW}$ 
11  | rebateScale = criticalLinks.size / noncriticalLinks.size
12  if rebateScale < 1 then
13    | rebate = rebate * rebateScale
14  end
15  foreach link  $l \in \mathcal{L}_c$  do
16    | # apply premium
17    |  $l.\textit{price} = l.\textit{cost} + (l.\textit{cost} * \textit{premium})$ 
18  end
19  foreach link  $l \in \mathcal{L}_n$  do
20    | # apply rebate
21    |  $l.\textit{price} = l.\textit{cost} - (l.\textit{cost} * \textit{rebate})$ 
22  end
23 end
```

Whereas WCRPA above attempts to pre-compute the cruciality of the bandwidth of particular links from a network-wide point of view, NBPPA performs direct computation of potential future consequences for each individual alternative course of action satisfying the individual traffic demand at hand. This is performed by measuring the change in the network's available resources if a given path is established on the network and calculating the effect the call has on all unique paths that share any of the selected path's links. The average blocking projection of the collection of feasible path alternatives determines whether the penalty the pricing algorithm applied comes in the form of a premium or a discount. Algorithm 2 describes how NBPPA applies pricing on paths.

Algorithm 2: Network Blocking Projection Pricing Algorithm

```

Input :List of feasible path alternatives  $\mathcal{P}$ 
1 averageBlockingProjection = getAvgBlockingProjection( $\mathcal{P}$ )
2 rebatePaths =  $\{ \forall p \in \mathcal{P} \mid p.\text{blockingProjection} <$ 
   averageBlockingProjection  $\}$ 
3 rebateScale = ( $\mathcal{P}.\text{size} - \text{rebatePaths}.\text{size}$ )/rebatePaths.size
4 # apply rebate/premium on the path's price
5 foreach  $p \in \mathcal{P}$  do
6   premium =  $|1 - (p.\text{blockingProjection} / \text{averageBlockingProjection})|$ 
7   delta = premium * p.originalPrice;
8   if  $p.\text{blockingProjection} < \text{averageBlockingProjection}$  then
9     # apply rebate
10    if rebateScale < 1 then
11      | delta = delta * rebateScale
12    end
13    p.price = p.originalPrice - delta;
14  else
15    # apply premium
16    p.price = p.originalPrice + delta;
17  end
18 end

```

5.2 Modeling Provider Resource Allocation Strategy

In the following chapters, we explore an ecosystem where network providers offer customers service choices, rather than only offering a single predetermined service. This choice-based network is based on the NSF funded research project ChoiceNet [RBCDGNW], which discussed the introduction on new architectural entities into the Internet, to enable economic interactions between stakeholders

and offer customers choices between services. The economic interactions would parallel real-world interactions that happens outside the network. ChoiceNet is reinforced by the principle that customers will reward well performing services, which will encourage competition between service providers to deliver more innovative services. Customers are enabled by a marketplace to discover, negotiate, and pay for services.

Innovative services made possible through flexible optical capabilities could resolve abstract service requirements for end-system or applications [TASPCP14]. Optical Packet and Circuit Integrated Network (OPCInet) is an example of an innovative solution for handling the service diversity that today's communication networks are faced with. OPCInet offers optical packet switching (OPS) links for bandwidth sharing applications and optical circuit switching (OCS) links for fully occupied bandwidth data transports with guaranteed end-to-end quality of service over a shared physical network [FSMHKSMTW12]. OPCInet allows a network operator to configure the network to dynamically switch links' wavebands (range of wavelengths) for either optical packet switching (OPS) and optical circuit switching (OCS) operation, enriching the network operators capability for handling a diverse set of traffic patterns. Interactive applications transmitting delay sensitive information such as distance-learning applications and video-conferencing or disaster recovering operations like bulk data transfers between data centers can benefit from such an architecture [VM14][VLZJ11]. A converged system that applies both the Choice-based network architectural approach with a flexible optical network would allow a provider to apply traffic grooming algorithms based on their current network state and advertise the most optimal set of service offerings for a customer to choose from [DRB13]. By advertising an optimal set of service offerings, a provider can efficiently use available optical resources and further improving performance of the optical network.

5.2.1 Design of OPCInet-inspired Study Scenarios

The OPCInet implementation's Wavelength-Division Multiplexing (WDM) links used a static spectrum assignment for OCS/OPS resources, with a set of flexibly assigned wavebands of shared spectrum for either OCS/OPS allocation based on operator's configurations [HFFMW14]. By statically assigning a separate waveband for OCS and OPS operation, the operator guarantees that at either extremes (majority OCS or OPS), a waveband exists for the minority resource along that link guaranteeing network-wide connectivity. The drawback from this approach occurs when underutilized (OCS/OPS) waveband exists, resulting in an increasing blocked/dropped rate for the opposite resource. Our implementation investigated a network where the entire spectrum of OCS/OPS wavebands are configurable. While OPCInet's resource adjustment algorithm performs waveband switching (OCS to OPS or OPS to OCS) after every circuit establishment [FMWH14][FSMHKSMTW12], our implementation instead performs this on a timer.

In our simulation studies, we characterize a specific study scenario by customer population,

Figure 5.2 Circuit Request Flow

one or more network providers with pricing algorithm (for each provider), a network topology, and a single provisioning strategy used for all providers network.

5.2.1.1 Interaction of Customers and Network Providers

The customer population, represented as a set of customer (user) requests, which arrive as either, an optical circuit or an optical packet. Each circuit request arrives from a unique choice-exercising customer with a pricing sensitivity limit before defection, while each packet request arrives from a single customer with no input on path selection. Packets are treated as best-effort traffic with no guarantee of delivery. The simulator demonstrates control over all the path pricing assignments and path alternative presentation for multiple providers. Each provider is equipped with a pricing algorithm and their set of network resources (set of switching OPCInet nodes and WDM links). Network-wide each link contained 16 wavelengths (λ)s assigned sequentially and continuous within wavebands (Λ) for every network provider. A waveband size of four wavelengths resolves to each link has 4 wavebands with (Λ_0) containing ($\lambda_0 - \lambda_3$) and (Λ_3) containing ($\lambda_{12} - \lambda_{15}$). Wavebands are allocated for either OCS or OPS operation and thus all wavelengths within the waveband's range are

*Waveband Configuration: Link has all OCS waveband(s) occupied or one unoccupied OCS waveband with at least one other OCS waveband's usage $\geq$ threshold, then leave allocation the same otherwise remove an OCS waveband

Figure 5.3 Resource Adjustment Algorithm

operated for that resource type. Our optical choice-based network study investigated the network configurations using the two extreme waveband sizes (one and four λ s) with an initial OCS:OPS allocation ratio of 1:3.

Customers request bandwidth between source and destination for either optical circuit or packet traffic; these requests are translated into a path, along a set of continuous wavelengths. Circuits may only be established using wavebands allocated for OCS operation and likewise, packets may only be successfully delivered if their path remains OPS operated throughout its transmission. Circuit requests are only accepted if the requested bandwidth is available throughout the path with a continuous set of wavelengths. Circuit requests may also be offered a Packet Waveband Sacrifice (PWS) enabled path alternative (PWSePA), a path that at the time of the request contains links with OPS allocated waveband(s), but that, if selected by the customer, would be immediately traded off (sacrificed) for one or more OCS wavebands. The selected waveband used during packet waveband sacrifice operation is determined by the provisioning strategy, discussed later in subsection 5.2.1.3. The circuit request flow is illustrated in Figure 5.2. Accepted circuit connections eventually depart

after their specified holding time has expired, freeing the wavelength for another circuit service or the waveband undergoes packet waveband reclamation by the resource adjustment algorithm, discussed in further details in the next subsection.

Packet requests on the other hand, are transmitted over an entire continuous waveband along a path to its destination [Har13], transmitting the packet with the sum of the throughput of all the wavelengths within the range of the waveband, this implies a direct correlation between packet throughput and the network's waveband size. The larger the waveband size, the faster the packet throughput, but fewer OCS/OPS combinations are produced. Packets are injected once at the source node on a waveband using round robin fashion, favoring the earliest occurring minimum occupied waveband's FDL [GTCT00]. Each link's waveband has an associated Fiber Delay Line (FDL). Fiber Delay Lines mitigate packet contention by offering packets different waiting times before servicing the packet. Packet contention occurs when two or more packets arrive at a switching node at the same time or the outgoing FDL is fully occupied. The current implementation of the FDL operates like a traditional packet-switch queue with a fixed size; an over estimation of how a true FDL behaves. Due to the changing nature of available packet waveband, packets can be dropped due to contention, during circuit waveband reclamation (OPS to OCS), or due to unavailable packet waveband at injection into the network. The packet traffic requirement of a continuous waveband for transmission, can also constrain packet transmission when OCS waveband is being underutilized. For example, if a circuit service occupied some wavelength (λ_0) then the entire waveband ($\lambda_0 - \lambda_3$, i.e. Λ_0) is handled by OCS. Despite there existing available wavelengths within this waveband ($\lambda_1 - \lambda_3$), the packet traffic's requirement of using an entire waveband combined with the strict waveband definition of being composed of a static set of continuous wavelengths, reduces the number of potential packet paths. The packet injection strategy attempts to minimize the number of dropped packets caused by such a dynamic system and increases packet bandwidth utilization.

Revenue is generated for every successfully established circuit and for every admitted packet. Packets successfully delivered are charged both an admission and delivery fees, while packets lost are only charged an admission fee.

5.2.1.2 Circuit Waveband Insertion and Removal Strategy

As discussed earlier, the our study's resource adjustment algorithm operates on a timer. During a resource adjustment algorithm run, the network performs circuit and packet waveband reclamation on all the network links. The adjustment algorithm switches a OPS waveband assignment to OCS (circuit waveband insertion), when the number of circuits in use reaches a threshold and moves OCS to OPS (circuit waveband removal) when there are no circuits in use within a waveband and the number of in use circuits falls below the threshold. A flow chart of its operation is illustrated in Figure 5.3. A circuit waveband insertion causes any occupied OPS waveband's FDL to drop any

stored packets before switching over to OCS. We have made circuit waveband insertion and removal separate configurable algorithms that can be implemented as strategies within the simulation. Together, the insertion and removal strategy defines the network's provisioning strategies, which are performed by the resource adjustment algorithm when several conditions are met.

OPCInet used a Single Boundary approach with OCS and OPS wavebands initially assigned together. With this approach waveband resource switching (OPS to OCS or vice-versa) moved along a shared boundary. For example, if the OCS wavebands are initially assigned the lower wavebands (Λ_{0-1}) and OPS wavebands are assigned the higher wavebands (Λ_{2-3}) then the highest indexed circuit waveband (Λ_1) will always border the lowest assigned OPS waveband (Λ_2). During circuit waveband insertion, OPCInet would select the lowest indexed (first-fit) OPS waveband (Λ_2) for OCS operation, while circuit waveband removal would select the highest indexed (last-fit) OCS waveband (Λ_1) for OPS operation as long as no active wavelength occupied that OCS waveband, otherwise the circuit waveband removal would not be permitted on the link. This systematic approach of moving resources, reduces the potential number of blocked circuit requests and dropped packets since all links with the same allocation ratios will always reserve the same wavebands and thus will have the same set of potentially available wavelengths. An alternative approach would be to allow any waveband to be selected for circuit insertion or removal based on a designated criteria, for example, a Free Placement approach that selects the OPS waveband with the fewest stored packets in their outgoing FDL during circuit waveband insertion and removes the highest indexed waveband among the unoccupied circuit waveband candidates during circuit waveband removal to reduce the spectrum fragmentation effect of this approach.

5.2.1.3 Packet Waveband Sacrifice enabled Path Alternative

The Packet Waveband Sacrifice (PWS) feature enables a provisioning strategy to preempt an active OPS waveband(s) to satisfy a circuit request. The resulting path alternative discovered using this mechanism is referred to as Packet Waveband Sacrifice enabled Path Alternative (PWSePA). Unlike the Resource Adjustment algorithm which operates on a timer, Packet Waveband Sacrifice occurs on a trigger and moves resources with a circuit requirement in mind. It performs this by pushing the boundary in only one direction, decreasing the packet share on a link at the cost of increasing the circuit share. Thus, Packet Waveband Sacrifice operates on infeasible path candidates in order to satisfy a circuit requirement, rather than on a per link basis like the circuit waveband insertion strategy. Additionally, our PWS implementation gives waveband preference to partially occupied circuit waveband paths before attempting to issue an unoccupied waveband per PWSePA discovery. The intention of this feature is to limit the additional spectrum fragmentation effects a PWSePA can have on the network. Figures 5.4 and 5.5 illustrates the flow of two PWS strategies studied. Each pricing algorithm persistently advertises a PWSePA candidate (based on the PWS implementation)

to every circuit request, if one is available.

PWS includes a penalty cost for the projected revenue loss from packet traffic due to realizing a circuit-service alternative. The projected revenue loss is based on the rate at which revenue is currently being generated by packets (carried and dropped) over a given link that would be sacrificed to realize the circuit-service alternative. This is summed over each link in the given circuit-service alternative and multiplied by the network's packet waveband reclamation estimate, the estimated time duration until this waveband is likely to be reclaimed for packets. The simulator determined this estimate by performing exponential smoothing on the network's average packet waveband reclamation duration for each link after each resource adjustment algorithm run.

In Chapter 9 we study the impact of using a Single Boundary strategy as opposed to the Free Placement strategy, and show that the latter is superior, as expected. All of the rest of our studies are based on the Free Placement approach for circuit waveband insertion/removal and uses the Free Placement PWS strategy.

*Each path may contain an unoccupied and/or occupied OCS waveband option. Path selection prioritize selecting paths with an occupied waveband to use as their selected waveband before offering paths with an unoccupied waveband.

Figure 5.4 Free Placement Packet Waveband Sacrifice Strategy

*Smallest Sequential Non-intersecting (SSN) Waveband: Next available OCS-capable waveband within a path's links dictated by the link(s) with the smallest OCS reservation

Figure 5.5 Single Boundary Packet Waveband Sacrifice Strategy

5.2.2 Pricing Algorithm Considerations for Hybrid Packet-Circuit Operation

Several considerations are needed on how our responsive pricing algorithms (described earlier) are to be implemented on an optical switch/router with capabilities for handling requests for both optical circuit and optical packet traffic on the same network links. Specifically, in this section, we define the metric for WCRPA's threshold, and make special consideration for NBPPA PWSePA's blocking projection.

WCRPA threshold is based on the link's OPS occupancy and thus a circuit waveband insertion/removal may shift the cost of a link, depending on whether the link's new OPS occupancy falls below the threshold.

Paths generated by PWS while using NBPPA can potentially grant network's blocking projection improvement by introducing additional circuit waveband(s) that could contain additional available circuit wavelength(s) for future circuit calls. This improvement can result in reducing the average

blocking projection and shift the balance of how NBPPA applies rebate/premiums to heavily favor the PWSePA paths with a rebate over premiums on the alternatives. This blocking projection improvement is not possible with a waveband size of one wavelength as the circuit service will always fully occupy the waveband and no available wavelength will be free from the PWSePA. To mitigate the effect that the PWSePA's blocking projection can have on its path alternatives' rebate/premium calculation, we have removed it from being included in the average blocking projection.

With the above considerations, the pricing strategies described earlier in Section 3.1 can be applied to the hybrid packet-circuit network.

5.3 Other Study Design Considerations

5.3.1 Traffic Arrival Traces

We refer to a list of connection requests, which can also be viewed as a history of connection requests, as a traffic arrival trace. The trace has a list of connection requests, each with an identifier. The identifier details when the connection request arrives in relation to the other requests. A trace consists of traffic uniformly distributed with any source-destination pairs equally possible. Circuit requests arrive along a Poisson-distribution and have a service time that is uniformly distributed between 1 and 5 minutes [VKC10].

As mentioned in our preliminary report, we previously considered creating traffic traces in which the magnitude of traffic was distance-dependent; that is, source-destination pairs separated by a longer distance across the network area have smaller traffic demands. Such traffic patterns have previously been considered realistic. However, we found that at this time, with the trend toward clouds and datacenters, such traffic patterns are generally not considered realistic; so we confined ourselves to traffic patterns with random variation in them, but no distance-dependent trend.

5.3.2 Path Alternative Discovery

Due to the random nature of customer requests, the network providers must perform dynamic routing and wavelength assignment. The state of the network changes over time as new circuit requests are established and released [Rou01]. For each circuit request, a network provider must check for available paths at the time of the request before offering a path in the form of an advertisement in the marketplace. To simplify providing an answer to the routing and wavelength assignment problem, each network provider offer path choices compiled using fixed-alternate routing algorithm [Jue01] for every source-destination node pairs. Fixed-alternate routing uses predetermined paths for establishing lightpaths. Alternate paths are composed using Yen's Algorithm for the k-shortest paths [Yen71][Smo17].

5.3.3 Numerical Results

A companion measurement tool was developed to automate the compilation of simulation results into graphical illustrations [Hun07]. Each graph gives the pricing algorithms measurements in context of its measured quantity. All comparisons between pricing algorithms are made using the same traffic arrival trace configurations. The following set of measured quantities are:

- Network Acceptance Rate: Average acceptance rate experienced by the network
- Circuit Noncompletion Rate: Average circuit noncompletion rate, circuit traffic not accepted by any network either because no path alternative was present and/or the customer's policy rejected all path alternatives
- Network's Average Link Utilization Running Average: Running average of each link's utilization
- Delay: Average accepted circuit path delay
- Total Revenue, Average: Total revenue from both traffic type (Circuit, Packet)
- Total Offered, Average: Total number of advertised Packet Sacrifice enabled routes

Each figure presents the statistical spread of the measurements in the form of 95% confidence intervals. A single provider figure reflects data using the same customer population but only a single implemented pricing algorithm. Measurements are categorized by the network's waveband size, with each colored vertical bar representing the measured quantity for a specific pricing algorithm, excluding outliers. On the other hand, a multi-provider scenario provides information about each provider categorized by their implemented pricing algorithm being presented along the X-axis, while each colored vertical bar indicates the experience of a different customer subpopulation for that provider (the bars represent the spread of customers from each subpopulation, excluding outliers).

5.3.4 Network Topologies

We picked the following common topologies 19-node European Optical Network (19EON), 24-node NSF Network (24NSF), and the 53-node Hibernia backbone network (53Hibernia). All of these topologies are well accepted in literature to receive realistic results and have been used by other researchers before us. Length of network links are based on the distances between switch nodes' coordinates. The following table breaks down the properties of each network topologies.

For one of our studies, we used a randomly generated topology to analyze the network performance of when multiple providers compete for business from the same customer population, and one has a smaller topology than the other. The graph we generated is based on the same nodal

Table 5.4 Simulation Parameters

Inter-arrival mean time: (Circuit = 750 ms, Packet train = 25 ms)
 Circuit Requested Bandwidth: 1 wavelength
 Circuit Customer Fractional Premium: .25, .5, .75, and 1
 Packet Train Size: 100 Mb, 1 Gb, 10 Gb
 Output FDL Size: 4 packets
 Each simulation has traffic arriving for 30 min
 Initial Optical Reservation: (Circuit = 4, Packet = 12)
 Circuit Waveband Occupancy Threshold: 70%
 Packet Waveband Reclamation Seeding Estimate: 90 seconds
 Exponential Smoothing Alpha: .5
 Exponential Smoothing Packet Waveband Reclamation Window Size: 20
 Exponential Smoothing Reclamation Window Size: 20
 Packet Waveband Injection: Round Robin with Minimum Occupied FDL
 Packet fee types: (Admission = 2, Delivery = 8)
 Packets are routed using the shortest path
 Yen Algorithm k : (Circuit customers = 15, Packet traffic = 3)

Table 5.5 Network Topologies

Name	Abbreviation	Description	Source	Nodal Coordinates	Figure
19-node European Optical Network	19EON	19 nodes with 39 edges	[FCTMJA99]	Table 5.6	5.6
24-node NSF Network	24NSF	24 nodes with 43 edges	[FDHTM15]	Table 5.7	5.7
53-node Hibernia backbone network	53Hibernia	53 nodes with 76 edges	[KNFBR11]	[Jam10]	5.8

Table 5.6 European Optical Network coordinates

ID	City Name	Coordinate
0	Oslo	59.9139, 10.7522
1	Copenhagen	55.6761, 12.5683
2	Stockholm	59.3293, 18.0686
3	Moscow	55.7558, 37.6173
4	Berlin	52.5200, 13.4050
5	Prague	50.0755, 14.4378
6	Vienna	48.2082, 16.3738
7	Zagreb	45.8150, 15.9819
8	Athens	37.9838, 23.7275
9	Milan	45.4642, 9.1900
10	Zurich	47.3769, 8.5417
11	Paris	48.85341, 2.3488
12	Madrid	40.4168, 3.7038
13	Lissa	51.8420, 16.5938
14	Dublin	53.34399, -6.26719
15	London	51.50853, -0.12574
16	Amsterdam	52.37403, 4.88969
17	Brussels	50.85045, 4.34878
18	Luxemburg	49.6116, 6.1319

Table 5.7 24-node NSF Network coordinates

ID	City Name	Coordinate
0	Seattle	47.60621, -122.33207
1	Portland	45.523064, -122.676483
2	Palo Alto	37.4419, -122.1430
3	Las Vegas	36.17497, -115.13722
4	San Diego	32.71533, -117.15726
5	Helena	46.5884, -112.0245
6	Salt Lake	40.7608, -111.8910
7	Tucson	32.2217, -110.9265
8	Boulder	40.0150, -105.2705
9	El Paso	31.7619, -106.4850
10	Chicago	41.85003, -87.65005
11	Lincoln	40.8258, -96.6852
12	Dallas	32.78306, -96.80667
13	San Antonio	29.4241, -98.4936
14	Columbus	39.9612, -82.9988
15	Frankfort	38.2009, -84.8733
16	Birmingham	33.5207, -86.8025
17	Mobile	30.6954, -88.0399
18	Ithaca	42.4440, -76.5019
19	Princeton	40.3573, -74.6672
20	Washington, DC	38.89511, -77.03637
21	Knoxville	35.9606, -83.9207
22	Raleigh	35.7721, -78.63861
23	Jacksonville	30.3322, -81.6557

coordinates as the 24NSFNet. Edges were randomly selected between nodes that are geographically close to each other. We define geographically close as being within 1/3 of the distance of the selected node's most distant node. By limiting our edges to a geographically close set, we achieve a more realistic topology and give the generated network a more competitive set of paths in comparison to the 24NSFNet. The 24-node generated graph (24Generated) and the consolidated graph between 24NSFNet and 24Generated topologies are illustrated in Figure 5.9 and Figure 5.9, respectively. Figure 5.10 displays the consolidated graph, with each distinct edge are color coded for a network provider (blue: 24NSFNet and red:24Generated) and common edges colored green.

Figure 5.6 European Optical Network topology (19-node European Optical Network with 39 edges)

Figure 5.7 NSF Network topology (24-node NSF Network with 43 edges)

Figure 5.8 Hibernia topology (53-node Hibernia backbone network with 76 edges)

Figure 5.9 24-node Generated network (24 nodes with 39 edges)

Figure 5.10 24-node network: NSF Network (blue edges) and Generated network (red edges) with mutual edges colored green

CHAPTER

6

IMPACT OF PRICING ALGORITHMS ON NETWORK PERFORMANCE

6.1 Research Question

In this chapter, we study what impact pricing algorithms have on a choice-based network's performance. To investigate this, we consider all other parameters and characteristics remaining the same in a scenario with a single choice-based network provider with a customer population, and examine the comparative network performance between scenarios in which the provider uses different pricing algorithms.

6.2 Experimental Setup

We explore each of the pricing algorithms discussed in Section 5.1 using a choice-based network scenario as described in Section 5.2.1. If pricing algorithms can have effect on the network performance, then pricing algorithms can be used to groom demand and assist in maximizing network utilization. Each pricing algorithm we studied was deployed on a 24-node NSF network topology (as detailed in Section 5.3.4) by a single network provider. Customer population distribution consisted of a 50:50 ratio of budget-limited customers with either a policy preference based on their delay-utility threshold or selecting the cheapest priced path. Simulation parameters are consistent with the

parameters mentioned in Table 5.4. Fifty simulations were run with each combination of pricing algorithm and waveband configuration.

Considering that PPA provides a baseline pricing algorithm, where no pricing adjustments were made on the offered path alternatives, we compare the network performance with PPA implemented against several different adaptive pricing algorithms to illustrate that pricing algorithms can have an impact on a choice-based network.

6.3 Experiment Results

Despite no significant difference being observed between each pricing algorithm's link utilization in Figure 6.1. Figure 6.2 illustrates the network experiencing different acceptance rates depending on the implemented pricing algorithm. The large waveband configuration allows for multiple wavelengths to become available during waveband readjustment. This new supply of multiple wavelengths favored responsive pricing algorithms' network state heuristic and allowed multiple wavelengths to be priced at the same rate.

Figures 6.3 and 6.4 show accepted customers experiencing different path delays per implemented pricing algorithm depending on their choice policy. Using PPA as the baseline, all responsive pricing algorithms, groomed their customers to using higher delayed paths on average. Maximum Utility customers' policy restrict their selections to paths with the lowest delayed path within their budget and thus will be the most inflexible to price change (i.e. if the path delay exceeds their threshold regardless of how low the price is, these customers will reject that path alternative). NBPPA is the only algorithm able to groom their traffic and achieve a similar acceptance rate as PPA.

In terms of revenue, Figure 6.5 shows that the studied pricing algorithms all generated significantly different values. The WCRPA based pricing strategy consistently generated more revenue than PPA. In all cases, this comes at the cost of accepting fewer customers and establishing longer delayed circuit paths when compared to PPA.

6.4 Conclusions

An extensive study was performed on a 24-node NSF network topology with several pricing algorithms. Data from this study shows definitively that pricing algorithms can have an impact on a choice-based network's performance. We have published our research contribution presented in this chapter in [UDMH17a].

Figure 6.1 Running Average Link Utilization

Figure 6.2 Circuit Acceptance Rate

Figure 6.3 Average Running Average Delay: Budget-limited Cheapest Available policy

Figure 6.4 Average Running Average Delay: Budget-limited Maximum Utility (with Utility Threshold) policy

Figure 6.5 Total Revenue Generated

IMPACT OF CUSTOMER CHOICE STRATEGIES ON NETWORK PERFORMANCE

7.1 Research Question

Collaborative optimizing of network resources relies on network providers adjusting cost on path alternatives based on their perceived scarcity and customers selecting optimal paths for their needs. If customer choices have no effect on network performance, then the interaction between customers and network providers cannot lead to maximizing network utilization. In this chapter, we investigate whether choice strategies have an impact on a choice-based network performance.

7.2 Experimental Setup

We explore three different proportion of customer populations on a multi-provider choice-based network ecosystem using the network simulator described in Section 5.2.1. In this study, there exist five network providers each using a unique pricing algorithm (as described in Section 5.1.2) on an identical network topology, specifically 24-node NSF network, with a network-wide waveband configuration of 4 λ (details of which were discussed in Section 5.3.4). Each network provider

competed with one another over a shared pool of customers, who only select paths based on their cost or its cost and delay. The customer proportion studied consisted of the proportions discussed in Section 5.1.1 and described in Table 5.3. Fifty simulations were ran with all parameters from Table 5.4 remaining the same but the customer distribution and the system's waveband size.

7.3 Experimental Results

We observe that among the five pricing algorithms, all experienced significantly different acceptance rates except PPA and WCRPA (75%), as seen in Figure 7.1. This can be explained by PPA inability to adjust advertised paths' cost thus being unable to steer customers toward underutilized paths, resulting in the provider's link utilization and acceptance rate having no significant difference regardless of the population's distribution. The same consequence occurred with WCRPA when set with a high threshold (75%), in this case, WCRPA high threshold aggressively penalized links with a higher premium as their circuit occupation increases, reducing the chance a customer will select a path that the provider views as using scarce resources. The link utilization performance is tied to the provider's acceptance rate and as result, Figure 7.2 matches the trend seen with Figure 7.1.

In respect to customers, all network providers delivered the lowest delayed paths to a customers with the 25:75 (cheapest priced and delay utility sensitive customers) proportion except for NBPPA, whose lowest delayed population came from the 75:25 (cheapest priced and delay utility sensitive customers) proportion (as seen in Figure 7.3). This disparity is a result of the difference with the customer's choice strategy. While both choice strategy are conscious of price (budget-limited), customers with delay-utility preference are far more restrictive then customers without, as can be seen in NBPPA's acceptance rate, which shows that far more circuit customers were accepted from proportion of 75:25 then 25:75. This trends highlights that the NBPPA provider's accepted paths were far more attractively priced and despite carrying a higher amount of delay. This results in NBPPA utilizing more network resources (links), thus explaining NBPPA high link utilization. On average, we observe in Figure 7.4 that NBPPA served Maximum Utility customers all experienced lower delayed paths per their QoS profile for every studied population when compared to PPA. Each QoS profile maintains a consistent trend influenced by their QoS Utility function, resulting in the Online Gaming based profile having the highest sensitivity to delay, while Voice Over IP based profile having the least sensitivity to delay.

All WCRPA-based providers earned more revenue with customers consisting with a higher proportion of delay utility sensitivity. This is expected because these customers would select a higher priced path alternative for its low delay despite their existing cheaper priced alternatives. For example this can be seen in Figure 7.5, where the customer composition consisting of 25:75 (cheapest priced and delay utility sensitive customers) generated more revenue for all three WCRPA-based providers. then the other two studied customer populations. As the customer proportion

decreases with the delay utility sensitive customers, so does the amount of generated revenue.

Figure 7.1 Circuit Acceptance Rate: Five service providers, each with an identical 24NSF topology

7.4 Conclusions

An exhaustive study on a multi-provider choice-based network with several different customer population proportion have revealed that customer choice does have an effect on a network's performance. Depending on the composition of the customer population, an informed service provider can provide themselves with a competitive advantage by selecting the right pricing algorithm that caters to their customer's preference. We have published our research contribution presented in this chapter in [UDMH17b].

Figure 7.2 Running Average Link Utilization: Five service providers, each with an identical 24NSF topology

Figure 7.3 Average Average Delay: Accepted Customer Population: Five service providers, each with an identical 24NSF topology

(a) 25% Budget-limited Cheapest Available policy and 75% Budget-limited Maximum Utility (with Utility Threshold) policy

(b) 50% Budget-limited Cheapest Available policy and 50% Budget-limited Maximum Utility (with Utility Threshold) policy

(c) 75% Budget-limited Cheapest Available policy and 25% Budget-limited Maximum Utility (with Utility Threshold) policy

Figure 7.4 Delay Experienced by Budget-limited Maximum Utility customers grouped by their QoS profile: Five service providers, each with an identical 24NSF topology

Figure 7.5 Total Revenue Generated: Five service providers, each with an identical 24NSF topology

CHAPTER

8

IMPACT OF WAVEBAND SIZE ON STRATEGY INTERACTION

8.1 Research Question

In this chapter, we repeat the study of the previous chapter with differing waveband sizes, to investigate how the granularity of packet/circuit provisioning affects the strategy interplay. Previously, we had focused on a waveband size of four wavelengths. In this chapter, we consider a “waveband size” of one wavelength; that is, every wavelength can be independently allocated to packet or circuit traffic, and independently switched.

8.2 Experimental Setup

For each waveband size, we again explore three different proportions of customer populations on a multi-provider choice-based network ecosystem using the network simulator described in Section 5.2.1. As before, there exist five network providers each using a unique pricing algorithm (as described in Section 5.1.2) on an identical network topology, specifically we study the competitive scenarios involving with 19-node European Optical network and 53-node Hibernia network (details of which were discussed in Section 5.3.4). Each network provider competed with one another over a shared pool of customers, who only select paths based on their cost or its cost and delay. The

customer proportions studied consisted of the proportions discussed in Section 5.1.1 and described in Table 5.3. Fifty simulations were ran with all parameters from Table 5.4 remaining the same but the customer distribution, the waveband size, and the network topology.

8.3 Experiment Results

8.3.1 Topology: 19-node European Optical network

We observe that NBPPA, our most sophisticated pricing algorithm, resulted in a win-win scenario for both customers and the provider when the network-wide waveband size is set at 4λ . In a win-win scenario both the service provider and the customers were able to collaboratively optimize the use network resources. For the service provider, this involves generating high network performance and while providing a customers with paths with low delay. NBPPA achieved this, by having the highest performing link utilization (Figure 8.1b), accepted the most customers (Figure 8.2b), and delivered the lowest delay for its customers (Figure 8.3b). NBPPA also generated the most revenue for all of the studied customer distributions but the 25:75 (cheapest available and maximum delay-utility). The loss in revenue is a result of the NBPPA adjusting favorable paths with low delay with higher premiums due to their projected impact on the network's blocking probability. Paths with lower delays are more desirable for delay-sensitive customers, and thus the impact NBPPA has on these customers are observed as the proportion of these customers increases. As a result of this price adjustment NBPPA can find itself unable to offer an acceptable path that is both within the customer's delay sensitivity threshold and under its budgetary constraint.

When we consider a network-wide waveband configuration consisting of waveband size of 1λ , NBPPA does not maintain its collaborative optimization status observed with the waveband size of 4λ . In fact, none of the network provider's pricing algorithm studied resulted in a win-win scenario with this configuration. While NBPPA continues to deliver the lowest delayed circuit paths to their customers (Figure 8.3a), it did not accept the most customers (Figure 8.2a). A single λ waveband reaches a congested state quicker then the same network using 4λ waveband because unused circuit wavebands are gradually converted to packet wavebands, once the network determines the circuit waveband is idle. A waveband greater then 1λ , is able to circumvent this conversion by having any one of the circuit wavelengths within that waveband in use. NBPPA traffic acceptance rate benefited from being able to reuse wavelengths within partially occupied wavebands with the 4λ set up, however, once that configuration was removed, NBPPA acceptance rate suffered resulting in WCRPA (75%) in accepting more customers. The higher acceptance rate results in WCRPA (75%) attaining a higher link utilization (Figure 8.1a) and the most revenue (Figure 8.4a).

8.3.2 Topology: 53-node Hibernia network

When we consider a larger topology where each of the five network providers advertised paths from identical 53-node Hibernia topologies, we observe that no significant difference between two highest link utilization among the providers using WCRPA (50%) and WCRPA (75%) when using the waveband size is 1λ as seen in Figure 8.5a. WCRPA (50%) accepted more traffic when serving the 75:25 (cheapest available : maximum delay-utility) population (Figure 8.6a) and delivered these customers with the lowest delayed paths (Figure 8.7a). Given these attributes, we see that network provider and customers are collaboratively optimizing their use of the network resources. WCRPA (75%) is able to also collaboratively optimize with the 25:75 (cheapest available: maximum delay-utility) population.

When a waveband size of 4λ is considered, WCRPA (50%) is observed collaboratively optimized with the customer population consisting of both 25:75 and 50:50 proportion of cheapest available and delay-utility customers. The low performance observed with the NBPPA, PPA, and WCRPA (25%) similarity between their network performance is a result of both the topology size and the simulation parameters not being as taxing as it was with the 19-node European Optical Network.

Overall waveband size had a significant effect on the pricing algorithm's competitive advantage in a multi-provider system. The waveband size factors into the network operator's decision of how to set the granularity of the packet-circuit sharing and will likely be reflected in costly hardware capabilities for the network nodes. Individual pricing algorithm showed competitive advantage against their competition, in terms of higher link utilization, higher customer acceptance, greater share of the revenue, or low average delay, depending on the provider's network topology and waveband size. But only with a waveband size of 4λ with a multi-provider environment involving 19EON topology and a waveband size of 1λ with the 53Hibernia topology, did a single pricing algorithm (NBPPA for 19EON and WCRPA (50%) for 53Hibernia) result in providing customers with the best experience on average, serviced the most customers, and took the largest share of the revenue among its competition; presenting this service provider and its customers with a win-win scenario. Thus we find that pricing algorithm selection and deciding on the granularity of packet-circuit sharing, are highly coupled.

8.4 Conclusions

Our results in this study showed that individual pricing algorithm can surface with competitive advantage over its competition in a multi-provider system offering services that use sophisticated optical capabilities such as hybrid packet-circuit switching irrespective of wavebanding granularity; however, the actual performances differ. This brings us to the important conclusion that the optimal choice of strategies cannot be considered a question that can be answered without reference to

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.1 Running Average Link Utilization: Five service providers, each with an identical 19EON topology

specific network engineering decisions such as waveband size. A provider must know its own network, as well as know its customers, to carefully set pricing policy in order to obtain the appropriate advantage from these capabilities. We have published our research contribution presented in this chapter in [UDMH18a].

(a) Network-wide waveband size of 1 λ

(b) Network-wide waveband size of 4 λ

Figure 8.2 Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology

(a) Network-wide waveband size of 1 λ

(b) Network-wide waveband size of 4 λ

Figure 8.3 Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.4 Total Revenue Generated: Five service providers, each with an identical 19EON topology

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.5 Running Average Link Utilization: Five service providers, each with an identical 53Hibernia topology

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.6 Circuit Acceptance Rate: Five service providers, each with an identical 53Hibernia topology

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.7 Accepted Customers Average Running Average Delay: Five service providers, each with an identical 53Hibernia topology

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 8.8 Total Revenue Generated: Five service providers, each with an identical 53Hibernia topology

IMPACT OF DIFFERENT RESOURCE PROVISIONING STRATEGIES IMPLEMENTED ON A MULTI-PROVIDER CHOICE-BASED NETWORK

9.1 Research Question

In this chapter, we consider the effect different resource provisioning strategies have on competing service providers' network performance for various customer populations. The resource provisioning strategy manages how a network of hybrid packet-circuit switching devices, implements their resource adjustment algorithm (circuit waveband insertion and removal) and determines what PWSePA to offer based on the network's current state. Specifically, we discuss the effect of two resource provisioning strategies, Free Placement and Single Boundary, described previously in Sections 5.2.1.2 and 5.2.1.3.

9.2 Experimental Setup

In this study, we analyze the results of a multi-provider choice-based network with two distinct resource provisioning approaches, Free Placement and Single Boundary. We explore these results with respect of the three customer population proportions discussed in Section 5.1.1 and described in Table 5.3. Five network providers compete for customers with each provider utilizing a distinct pricing algorithm. We will focus on two competitive environments, distinct by their topologies: the 53-node Hibernia network and the 19-node European Optical network. We initially discuss the different effect each provisioning strategy have on the system's blocking rate, for competitive environments with several different network topologies and follow that with an in-depth comparative analysis on a Single Boundary provisioned environment. Fifty simulations were ran with all the parameters from Table 5.4 remaining the same but the customer distribution and the system's waveband size.

9.3 Experiment Results

The strict methodology of Single Boundary restricts the number of allowable circuit waveband insertion/removal operations and PWSePA offerings. By looking at the competitive environment's blocking rate, we are informed on the efficiency of the combined force of the multi-provider's pricing algorithms in accepting calls from a shared pool of customers. By comparing the results of different provisioning strategies with respect to their shared waveband size, we can study the cumulative effect the provisioning strategy methodology had on the system's blocking rate. As a system, Figures 9.1 and 9.2 show that fewer circuits were established, with a Single Boundary multi-provider environment when the network-wide waveband size of 1λ which can be observed for every customer proportion studied. But when a waveband size of 4λ is considered, only the competitive environment with service providers' offering paths from a 53-node Hibernia backbone topology experienced fewer blocked calls with Free Placement provisioning, otherwise the system experienced significantly fewer blocked calls with Single Boundary. In the following subsections, we will provide an in-depth discussion on these results for the two competitive environments.

9.3.1 Provisioning strategy effect on 53-node topology

The 53-node topology competitive environment's blocking rate increased as much as 700% and 300% with a 1λ and 4λ waveband network, respectively, in comparison to their Free Placement's values regardless of the population they served. The large increase in blocked calls also hints that a number of circuit requests were being refused service due to a lack of path resources at the time of the request. Insufficient path resources is a result of either, the network being heavily utilized or an inefficient resource (circuit/packet waveband switching) allocation. Figure 9.3 points to an

inefficient allocation of resources being the cause of the high blocking rate, due to Single Boundary providers experiencing a lower link utilization in comparison to their Free Placement's competitive environment's utilization (Figures 8.5). PWSePA offers network providers the ability to remain competitive in dynamic situations when the network has inadequate amount of resources to satisfy a path request due to the network's current resource allocation arrangement. Single Boundary significantly dropped the number of the PWSePA offered, regardless of the system's waveband size being considered (Figures 9.4 and 9.5).

Figure 9.1 Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 53Hibernia topology

9.3.2 Provisioning strategy effect on 19-node topology

We now consider the specific system where we observed both service providers and customers collaboratively optimizing the network resources in Section 8. Within this environment, five providers offered path services from an identical 19-node European Optical network. Performance metrics of the Single Boundary provisioning strategy are illustrated in Figures 9.6 - 9.9.

As we mentioned earlier in this section, the system's blocking rate (Figure 9.2) improved with the Single Boundary provisioning strategy with a waveband size of 4λ . This is the result of an efficient allocation of circuit waveband resource, based on the fact that the providers' link utilization with the two provisioning strategies remained roughly the same, with respect to the pricing algorithm. As expected, regardless of the networks' waveband size, Single Boundary approach significantly

Figure 9.2 Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 19EON topology

(a) Network-wide waveband size of 1 λ

(b) Network-wide waveband size of 4 λ

Figure 9.3 Running Average Link Utilization: Five service providers, each with an identical 53Hibernia topology using Single Boundary Provisioning (compare with Figure 8.5)

(a) Single Boundary strategy

(b) Free Placement strategy

Figure 9.4 Total PWSePA offered: Five service providers, each with an identical 53Hibernia topology with waveband size of 1λ

(a) Single Boundary strategy

(b) Free Placement strategy

Figure 9.5 Total PWSePA offered: Five service providers, each with an identical 53Hibernia topology with waveband size of 4λ

reduced the average number of PWSePA offered when compared to the Free Placement approach (Figures 9.10 and 9.11), following a similar trend witnessed in the previous subsection with a larger topology.

Specifically looking at the effect that the Single Boundary approach has on the service provider that resulted in win-win situation, it can be observed that the gentle systematic approach of this methodology improved NBPPA performance in terms of its customer acceptance rate and revenue earned with a waveband size is 1λ , when compared to similar results when Free Placement provisioning is implemented (Figures 8.1 - 8.4). This improvement resulted in NBPPA accepting a larger share of customers when the population proportion consisted of 75:25 and 50:50 (cheapest available:maximum-utility) then Free Placement's higher acceptance rate provider (WCRPA-75).

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 9.6 Running Average Link Utilization: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning (compare with Figure 8.1)

9.4 Conclusions

We have discussed in details the impact of two resource provisioning strategies within several competitive scenarios involving multiple service providers with various network topologies and customer populations. Our data show that the methodology of the resource provisioning strategy have a large impact on the efficiency of the system. Based on these results, a service provider should carefully consider their implemented provisioning strategy based on their network topology, waveband setting, and expected customer policies.

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 9.7 Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 9.8 Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning

(a) Network-wide waveband size of 1λ

(b) Network-wide waveband size of 4λ

Figure 9.9 Total Revenue Generated: Five service providers, each with an identical 19EON topology using Single Boundary Provisioning

(a) Single Boundary strategy

(b) Free Placement strategy

Figure 9.10 Total PWSePA offered: Five service providers, each with an identical 19EON topology with waveband size of 1λ

(a) Single Boundary strategy

(b) Free Placement strategy

Figure 9.11 Total PWSePA offered: Five service providers, each with an identical 19EON topology with waveband size of 4λ

CHAPTER

10

IMPACT OF NETWORK PROVIDERS OFFERING CUSTOMERS A SINGLE PREDETERMINED PATH CHOICE

In this chapter, we examine the impact of providers offering only a single predetermined path within a multi-provider environment where each service provider compete for business from a single pool of customers. The predetermined path is determined based on the provider's path provisioning strategy and limits customers' choice to a single advertised path. This work expands on earlier work discussed in Chapters 7 and 8, which looked at the same problem for a single service provider and across various path provisioning strategies. By studying the effect of each path provisioning approach has on the network, we can determine the effectiveness customer choice has on the network.

10.1 Experimental Setup

In this study, we explore this problem with the 19-node EON network (described in Section 5.3.4) and analyze the results of several simulation experiments consisting of network-wide path provisioning strategy being implemented by the studied service providers. In each of these experiments, the competitive landscape involves a service provider utilizing one of the pricing algorithms discussed

in Section 5.1.2 and a pool of customers. The customer pool studied consist of the distributed set of customer strategies discussed in details in Section 5.1.1. Our results are based on aggregated data compiled from fifty simulation runs for each combination of path provisioning strategies and customer population distributions. The remaining chapter follows the following design: brief background of the studied path provisioning strategies, a follow-up on system's blocking rate for each of the studied path provisioning strategies, and finally a deep dive analysis into the performance of the non-choice path provisioning strategies on their network.

10.2 Background: Path Provisioning Strategies

In Chapter 7, we discussed the effect of customer choice has on the performance of multiple providers' network. With respect to path provisioning, we would regard this setup as having choice path provisioning strategy, since the providers offered multiple available path alternatives for the customer to choose from. A path provisioning strategy that only offers a single predetermined path would be regarded as being a non-choice provisioning strategy. With non-choice provisioning strategies, a path is selected based on a specific criteria: cheapest path (customer's most precious path, in terms of cost), shortest delayed path (customer's most precious path, in terms of delay), or the path with the highest profit margin (provider's most profitable path). The profit-oriented path provisioning approach selects the path with the largest profit margin determined by the difference between the pricing algorithm's adjusted price and the path's base price. In the case of a tie, the non-choice provisioning strategy selects a path at random.

10.3 Results and Discussion

With respect of the path provisioning strategies, both the shortest-delay and cheapest path are customer-centered approaches that view higher utilized resources can only be achieved by eliminating the paths that utilize the fewest amount of resources or the fewest scarce resources, respectively. While the most profitable path provisioning strategy is a provider-centered approach where profit margins are used to give an outlook on what paths will lead to greater revenue and utilized resources. In this section we will discuss the effect each of these path provisioning strategies had on a competitive network system's overall blocking rate and follow up this with an analysis on a broader look into the cheapest path provisioning strategy had on individual competing service providers for three customer population proportion.

Results show in Figure 10.1 that no significant difference can be observed in the circuit noncompletion rate between choice and two non-choice path provisioning strategies: cheapest available path and shortest delay, regardless of the network's waveband size. The close relationship between the two non-choice strategies can be best explained by looking at the underlying metric each

customer choice strategy preferences are based on and the applied path provisioning strategy. Delay-utility policy favors the shortest available distances, while cheapest available policy resolves to the same path selected by the cheapest available path. With respect to the system's waveband size, none of populations studied with the most profitable path provisioning strategy experienced any significant difference in their blocking rate.

Figure 10.1 Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 19EON topology

In Chapter 8, we discovered that a win-win situation can occur for competitor using the Network Blocking Projection pricing algorithm (NBPPA) within a competitive environment involving choice services offered from a 19-node EON topology with a network-wide waveband size of 4λ implemented. Despite the overwhelming similarities between choice and the two non-choice strategies discussed earlier, none of the studied pricing algorithms experienced a competitive win scenario for both customers (in terms of experienced delay) and a provider (in terms of revenue and link utilization). In a matter of fact, NBPPA consistently under-performed against its peers with all measured metrics except link utilization (which it remained comparable to), regardless of the waveband size, as can be seen in Figures 10.2 - 10.5. For example, NBPPA received nearly 300% more of customer population with a 75:25 (cheapest available and maximum delay-utility) proportion when a Choice

paradigm is enabled (Figure 8.2b compared to Figure 10.4). The lack of path options limits NBPPA ability to balance out the provider's network performance based on offered path services. This leads us to the conclusion that the competitive advantage witnessed with NBPPA with Choice paradigm is the result from the collaborative competitive system where service options are transparently being offered and allowed to be diversely selected.

In regards of the network performance between the two non-choice path provisioning strategies, we can see that despite the greedy profit-oriented ideology of the Most Profitable path provisioning strategy, their providers generated far less revenue then their peers using the Cheapest Available Path provisioning strategy. Among the studied performance metrics, Most Profitable path provisioning strategy providers experienced higher link utilization, but served fewer customers and offered those who did received services a higher variance of delay.

(a) Cheapest Available Path Provisioning

(b) Most Profitable Path Provisioning

Figure 10.2 Running Average Link Utilization: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figure 8.1b)

10.4 Conclusions

In this chapter, we studied the implication of withdrawing service choices from a multi-provider marketplace and found that with all things remaining the same except for the system's path provisioning strategy, this approach resulted in a drastically different ecosystem compared to the choice-enabled implementation, in terms of the customers experience and the networks' performance and uti-

(a) Cheapest Available Path Provisioning

(b) Most Profitable Path Provisioning

Figure 10.3 Accepted Customers Average Running Average Delay: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.3b)

(a) Cheapest Available Path Provisioning

(b) Most Profitable Path Provisioning

Figure 10.4 Circuit Acceptance Rate: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.2b)

(a) Cheapest Available Path Provisioning

(b) Most Profitable Path Provisioning

Figure 10.5 Total Revenue Generated: Five service providers, each with an identical 19EON topology with a different non-choice path provisioning and network-wide waveband size of 4 (compare with Figures 8.4b)

lization, in both of the network topologies that we studied. While this study looked into several non-choice path provisioning strategies, none of the studied strategies resulted in both the customer population and a provider experiencing a win-win situation, unlike the one witnessed with choice-enabled path provisioning. We have studied a similar setup with a single network provider in [UDMH17b] and came to same conclusion.

CHAPTER

11

IMPACT OF EXCHANGE POINTS IN A MULTI-PROVIDER CHOICE BASED NETWORK

In this chapter, we discuss the effect of introducing exchange points, similar to Internet Exchange Points (IXPs), into a marketplace of network services, where multiple service providers offer uniquely priced services to satisfy traffic demands. The establishment of exchange points allows network providers to cooperatively offer path services that satisfies a customer's request while reducing the impact on provider's network. This cooperation leads to an increase in the number of path services the customer population can choose from and provide a reduced utilization impact on each provider's network. Measuring the impact of cooperating network providers can provide network managers with ideas on which competing provider they should consider partnering with in order to achieve higher revenue while reducing network maintenance.

11.1 Experimental Setup

In this study, we explore this problem with the 24-node NSF network (described in Section 5.3.4) and analyze the results of several simulation experiments consisting of network-wide path provisioning strategy being implemented by the studied service providers. In each of these experiments, the

competitive landscape involves five service providers each utilizing one of the pricing algorithms discussed in Section 5.1.2 and a pool of customers. Each service providers have access to offering path alternatives through a set of shared exchange points (Nodes 2, 10, 12, and 20). The customer pool studied consist of the distributed set of customer strategies discussed in details in Section 5.1.1. Our results are based on aggregated data compiled from fifty simulation runs for each combination of path provisioning strategies and customer population distributions.

11.2 Experimental Results

The following results will look at the effect of including exchange points into multi-provider choice-based network. First by comparing the impact of adding exchange nodes into the marketplace by comparing system with exchange nodes against one without. Following this, we will look at which pricing algorithms faced a disadvantage/advantage as a result of enabling exchange points within a competitive environment.

11.2.1 Effect of enabling exchange points

With cooperation between network providers in offering path services, we would expect to see more circuits accepted because the service providers are sharing a mutual burden of carrying traffic. Instead we find that with all things remaining the same, except for the exchange nodes being enabled, both competitive scenarios resulted in a similar trends for the studied populations. Both setups experienced similar circuit noncompletion rates (Figure 11.1), with a slight variance appearing to show that the inclusion of more exchange points could result in a reducing the circuit noncompletion rate further for the exchange point system. There is no significant difference observed for the service providers, in terms of link utilization (Figure 11.4) despite the configured waveband configuration. NBPPA benefited the most from not including the exchange points when the waveband size was 4λ , in terms of circuit establishments, (Figure 11.3 compared to Figure 11.2) with all the studied customer populations. With respect to this study, only the originating network provider gets credit for the circuit establishment, and thus the reduction in accepted traffic seen with the exchange point scenario combined with matching link utilization (Figure 11.4 compared to Figure 11.5) hints that the provider's resources are remaining occupied with competition's traffic. Paths offered over a competitor's network can involve using their scarcer resources, (which that competitor would have priced as such) increasing the chance of that competitor not being able to fulfill future requests. This explains why the acceptance rate for NBPPA reduced when exchange points were included in the study, the other competitors are benefiting from sending traffic across NBPPA's network.

In terms of revenue, all of the providers earned less when the exchange points were utilized (Figure 11.6 compared to Figure 11.7). This is a side effect of service providers co-operating to carry

traffic. Accepted traffic carried over composed paths are the result of utilizing the cheapest resources from multiple providers and can result in being among the shortest paths (maximum utility) being offered. As a result, all providers end up losing revenue which would have been generated from satisfying the request themselves by using premium (scarce) resources. This is observed the most when comparing the system's with a waveband size was 4λ . A 1λ waveband network, has fewer opportunities to benefit co-operating service providers because of the fragmentation of the waveband traffic (circuit/packet). Customers benefit from these composed paths, which can be seen with all the studied populations and providers, when comparing the average delay a customer experiences in Figures 11.8 against 11.9.

In regards to PWSePA, a feature far less required for the dynamic pricing algorithms studied because of their ability to adjust their price as a function to the current congestion faced on their network. These providers are able to create a load balancing distribution of traffic with their pricing algorithm, resulting in less of a need to offer PWSePA. Traffic being the same, a rigid resource pricing algorithm like PPA is going to reach that point earlier for congested parts of its network in comparison to the other pricing algorithms studied. The dynamic pricing algorithms ability to adjust pricing with congestion will only reach this point once congestion is uniformly distributed across it's network. Enabling exchange points extends this point from being reached even further because providers are now able establish circuits across co-operating competition's network. This can be seen in the reduction in PWSePA being accepted in Figures 11.10 and 11.11. This is further supported when considering that customers who selected PPA experienced the highest amount of delay, as observed in Figure 11.8.

As the proportion of customers with preferences for selecting the cheapest priced paths increase, PPA is seen requiring more PWSePA. This is the result of congestion on the PPA network, being less likely to offer an attractive path that is both within the customers' budgetary constraint and priced cheaper than its competition. This pattern holds true regardless of the network waveband deployed within the competitive scenario.

11.2.2 Disadvantaged Pricing Algorithm

With exchange points enabled, competition between the service providers demonstrated a clear disadvantage for the provider using WCRPA-25, in terms of circuit establishment rate regardless of the population proportion or network waveband configuration as illustrated in Figure 11.2b. The low circuit establishment resulted in lower link utilization and total revenue (as seen in Figures 11.4 and 11.6), this is to be expected since link utilization measures network link circuit occupancy and revenue is primarily generated from circuit traffic. In most of the WCRPA-based providers, as the providers competed for a larger proportion of customers with a preference for the cheapest priced path, these providers performed worse, in terms of establishing fewer circuits, lower link

utilization, and less revenue. The economical forces of both the pricing adjustment of WCRPA and these customers' sensitivity to the pricing of paths, dissuaded these customers to accept a better priced offering from competition.

11.2.3 Advantaged Pricing Algorithm

Depending on the network waveband deployed on the network and the population proportion being served, a different pricing algorithm appears to have an advantage over the competition. With a network wide waveband of 1λ , no single provider is able to monopolize in both link utilization, revenue, and circuit establishment. WCRPA-75 is able to establish a larger percentage of circuits despite the populations' choice strategies proportion but only commanded a larger amount of revenue with the 25:75 (cheapest available preference, delay sensitive preference) proportion. There is no significant difference among the top competitor's link utilization. When the competition involved a network wide waveband of 4λ , NBPPA has the competitive edge against its competitors achieving an higher link utilization and accepting a larger portion of circuit traffic for all the studied population proportion. This is not unexpected, as link utilization corresponds with the amount of circuit traffic transmitting through the service provider's network. With respect to revenue, NBPPA achieves an equitable amount of revenue in range of other top competitors with no significant difference being observed among them. NBPPA ability to load balance traffic across its network, provides its customers a larger supply of lower delayed paths across its network in comparison to WCRPA-based providers. This can be observed in the experienced delay which the customer population faced, when they selected NBPPA as its service provider, as seen in Figure 11.8.

11.3 Conclusions

The distinction between the pricing algorithms we studied shows that a pricing algorithm does make a difference on what market share of the population a service provider is able to attain. No single pricing algorithm was a clear winner, without considering the different distributions of customer preferences and the configuration of the network; thus a service provider must choose their pricing algorithm with consideration of these factor in order to provide themselves with a competitive edge. Future studies should consider investigating different exchange node placements and network topologies to further understand the effect pricing algorithms have on the network economics on the system.

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.1 Running Average Circuit Noncompletion Rate: Five service providers, each with an identical 24NSF topology and a different pricing algorithm

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.2 Successful Circuit Demands: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.3 Successful Circuit Demands: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.4 Running Average Link Utilization: Five service providers, each with an identical 24NSF topology with an different path provisioning strategy while offering path alternatives using exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.5 Running Average Link Utilization: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.6 Total Revenue: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.7 Total Revenue: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.8 Experienced Delay: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.9 Experienced Delay: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.10 PWSePA Successfully Offered: Five service providers, each with an identical 24NSF topology and a different pricing algorithm. All offering path alternatives that may use exchange points

(a) Network-wide waveband size 1λ

(b) Network-wide waveband size 4λ

Figure 11.11 PWSePA Successfully Offered: Five service providers, each with an identical 24NSF topology and a different pricing algorithm but without exchange points

IMPACT OF COMPETITION BETWEEN MULTIPLE PROVIDERS WITH DIFFERENT TOPOLOGIES

In this chapter, we discuss a multi-provider scenario consisting of two competing service providers offering network paths for a shared pool of customers. Each topology contain the same set of co-located network nodes, but the two possess different sets of fiber connections between their network nodes.

12.1 Experimental Setup

We explore the impact of imbalanced competition between two competing choice-based network providers, by analyzing the effect of the pricing algorithms discussed in Section 5.1 deployed on two unique networks using co-located network nodes. We specifically look at the 24-node networks labeled 24NSFNet and 24Generated in Section 5.3.4. The combined topology view of these two networks is illustrated in Figure 5.10. These networks will compete over a population of customers deploying 25:75 proportion of cheapest available : maximum utility strategies. This is one of the population proportion discussed in detail in Section 5.1.1 and described in Table 5.3.

12.2 Experimental Results

We will look at the impact of these competitive scenarios from two views, when the same pricing algorithms are being implemented by the two networks and another when the pricing algorithms are different.

12.2.1 Competitive Scenario: Same Pricing Algorithms

In this section, we explore the effects of two service providers competing over a mutual population of users while deploying the identical pricing algorithms. From Figure 12.1 we observe that Network A consistently received a larger proportion of circuit customers, regardless of the network-wide waveband size being studied, which resulted in Network A consuming a higher link utilization (Figure 12.3). As a consequence to receiving less circuit traffic, the economic forces of the system caused Network B's waveband allocation to favor packet switching. This is evident when considering that the customers who did select Network B benefitted from shorter delays, which can be observed in Figure 12.2, regardless of network-wide waveband implemented. As expected, the service provider (Network A) which maintained a higher link utilization and accepted the more circuit traffic generated the largest amount of revenue (see Figure 12.5).

When we consider PWSePA, we observe that PWSePA is less beneficial for Network A when it is using the 4λ waveband configuration, because this waveband configuration allows the provider to accept future traffic using the residual wavelengths from the waveband after utilizing PWSePA. The opposite is true for a 1λ waveband configuration, where future traffic can not benefit from providing a PWSePA. This is illustrated in Figure 12.4. We conjecture that if waveband configuration continues to be increased, we would observe PWSePA to continue to trend toward offering the provider no beneficial attributes.

12.2.2 Competitive Scenario: Different Pricing Algorithms

In this section, we explore the effects of two service providers competing over a mutual population of users while deploying different pricing algorithms. Based on the studied results, none of the competitive scenarios achieved a win-win situation between the two network providers. In all competitive scenarios, one network provider (Network A) maintained a competitive edge over the other in terms of circuit establishments and revenue (Figures 12.6, 12.7, 12.8 and 12.9), while the customers served by the under achieving network provider (Network B) experienced lower delayed paths in all scenarios except when competing against Network A's with NBPPA deployed and a network-wide waveband size of 4λ (Figure 12.11b). There is little to no significant difference observed between pricing algorithms for Network A in terms of circuit acceptance rate and total revenue, with respect to the waveband configuration. Network B achieved its best (acceptance rate)

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.1 Successful Circuit Establishment: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.2 Experienced Delay: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.3 Running Average Link Utilization: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.4 Portion of Selected Packet Waveband Sacrificed enabled Path Alternative: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.5 Total Revenue: Two service providers, each with an identical Pricing Algorithms and network-wide waveband sizes but different topologies

results when competing against Network A's WCRPA-75. For Network A, the pricing algorithm that provided the best circuit acceptance rate (WCRPA-25) did not generate the most revenue (WCRPA-50), yet Network A's customers received a better experience using WCRPA-25 with network waveband size of 1λ and showed no significant difference against WCRPA-50 with a network waveband size of 4λ .

The asymmetric dominance witnessed in this study between the two service providers regardless the network wide waveband configuration or the pricing algorithm deployed, suggests that for the given simulation parameters Table 5.4, Network A is as similar to Network B as expected.

12.3 Conclusions

In the chapter, we studied a multi-provider scenario consisting of two competing service providers offering nearly similar network paths for a shared pool of customers and found that with all things remaining the same except the provider's pricing algorithms, the observed system resulted in one network provider consistently dominating the other, in terms of network performance and link utilization. While this study looked into several competing scenarios, none of the studied scenarios resulted in both the customer population and either providers experiencing a win-win situation. Overall, this was a surprising finding, and we conjecture that the limited size of the network topologies may have caused the effect of small differences between the two to have a disproportionate effect in advantaging one over the other. Further investigation into these pricing

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.6 Successful Circuit Establishment: Two service providers, each with a different Pricing Algorithm and different topology

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.7 Successful Circuit Establishment: Two service providers, each with a different Pricing Algorithm and different topology

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.8 Total Revenue: Two service providers, each with a different Pricing Algorithm and different topology

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.9 Total Revenue: Two service providers, each with a different Pricing Algorithm and different topology

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.10 Experienced Delay on Network A when competing against another service provider, each with a different Pricing Algorithm and different topology

(a) Network Waveband size of 1λ

(b) Network Waveband size of 4λ

Figure 12.11 Experienced Delay on Network B when competing against another service provider, each with a different Pricing Algorithm and different topology

algorithms with larger topologies would expand the understanding how these pricing algorithms react with larger or more dense networks.

CHAPTER

13

CONCLUSIONS

In this thesis, we embarked on a study of how provider and customer strategies interact in a choice-based network scenario, and whether such interaction can be mutually beneficial for the right choice of strategies.

Our overall conclusion is that such strategies can indeed have an effect, but the nature and magnitude of the beneficial effect is sensitive to several factors, including network capabilities such as wavebanding granularity and packet-sacrifice to open up circuit wavebands, and customer population characteristics such as the mix of customers with different utility functions.

While we undertook a thorough study with a specific reasonable set of such capabilities and customer populations, our very conclusions show that such a study must be repeated for specific scenarios with particular sets of network capabilities. As the demand for optical transport continues to grow without bound in the planetary network, and new technological capabilities in optical communication and networking continue to be deployed in response, such scenarios will continually evolve into new ones. We hope that our research in this thesis and our publications inform such future studies, both by academic researchers, and practitioners in their architecture and design of networks utilizing ever-newer device and network capabilities.

BIBLIOGRAPHY

- [Ari13] Aricent. "Application-Aware Routing in Software-Defined Networks". 2013. URL: http://www.aricent.com/pdf/Aricent_Whitepaper_-_Application_Aware_Routing_in_SDN.pdf.
- [BD13] Babaoglu, A. & Dutta, R. "A Verification Service Architecture for the Future Internet". *Computer Communications and Networks, 2013. ICCCN 2013. Proceedings. 22nd International Conference on*. 2013.
- [BD14a] Babaoglu, A. & Dutta, R. "A GENI Meso-Scale Experiment of a Verification Service". *Research and Educational Experiment Workshop (GREE), 2014 Third GENI*. 2014, pp. 65–68.
- [BD14b] Babaoglu, A. & Dutta, R. "Grooming demands instead of traffic: Benefits of a choice-based approach to Green optical networks". *Global Communications Conference (GLOBECOM), 2014 IEEE*. 2014, pp. 2029–2035.
- [BD14c] Babaoglu, A. C. & Dutta, R. "Grooming Demands instead of Traffic: Benefits of a Choice-Based Approach to Green Optical Networks". *Proceedings of IEEE Globecom, Austin, TX*. 2014.
- [BUDR16] Bhat, S., Udechukwu, R., Dutta, R. & Rouskas, G. N. "Inception to application: A GENI based prototype of an open Marketplace for network services". *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. 2016, pp. 1043–1044.
- [BUDR17a] Bhat, S., Udechukwu, R., Dutta, R. & Rouskas, G. N. "Network service orchestration in heterogeneous 5G networks using an open marketplace". *IET Networks* **6.6** (2017), pp. 149–156.
- [BUDR17b] Bhat, S., Udechukwu, R., Dutta, R. & Rouskas, G. N. "On service composition algorithm for Open Marketplaces of network services". *2017 European Conference on Networks and Communications (EuCNC)*. IEEE. 2017, pp. 1–6.
- [BD20] Brockelsby, W. & Dutta, R. "Performance Implications of Problem Decomposition Approaches for SDN Pipelines". *GLOBECOM 2020-2020 IEEE Global Communications Conference*. IEEE. 2020, pp. 1–6.
- [Chi12] Chiosi, M. et al. "Network Functions Virtualisation". White paper. 2012. URL: portal.etsi.org/NFV/NFV_White_Paper.pdf.
- [DRB13] Dutta, R., Rouskas, G. & Baldine, I. "Converging Choice and Service in Future Commodity Optical Networks using Traffic Grooming". *15th International Conference on Transparent Optical Networks (ICTON 2013)*. 2013, pp. 1–5.

- [Ell08] Elliott, C. "GENI - global environment for network innovations". *Local Computer Networks, 2008. LCN 2008. 33rd IEEE Conference on*. 2008, pp. 8–8.
- [FVA06] Farrel, A., Vasseur, J.-P. & Ash, J. *A Path Computation Element (PCE)-Based Architecture*. RFC 4655. 2006. URL: <http://www.ietf.org/rfc/rfc4655.txt>.
- [FDHTM15] Ferdousi, S., Dikbiyik, F., Habib, M. F., Tornatore, M. & Mukherjee, B. "Disaster-Aware Datacenter Placement and Dynamic Content Management in Cloud Networks". *J. Opt. Commun. Netw.* **7.7** (2015), pp. 681–694. URL: <http://jocn.osa.org/abstract.cfm?URI=jocn-7-7-681>.
- [Fin12] Finnie, G. "The Role of DPI in an SDN World". *Heavy Reading*. Qosmos, 2012. URL: http://www.qosmos.com/wp-content/uploads/2013/03/Heavy_Reading-Qosmos-DPI-SDN-WP-Dec-2012.pdf.
- [FCTMJA99] Fumagalli, A., Cerutti, I., Tacca, M., Masetti, F., Jagannathan, R. & Alagar, S. "Survivable networks based on optimal routing and WDM self-healing rings". *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*. Vol. 2. 1999, 726–733 vol.2.
- [FMWH14] Furukawa, H., Miyazawa, T., Wada, N. & Harai, H. "Moving the boundary between wavelength resources in optical packet and circuit integrated ring network". *Opt. Express* **22.1** (2014), pp. 47–54. URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-22-1-47>.
- [FSMHKSMTW12] Furukawa, H., Shinada, S., Miyazawa, T., Harai, H., Kawasaki, W., Saito, T., Matsunaga, K., Toyozumi, T. & Wada, N. "A Multi-Ring Optical Packet and Circuit Integrated Network with Optical Buffering". *Opt. Express* **20.27** (2012), pp. 28764–28771. URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-20-27-28764>.
- [GTCT00] Ge, A., Tancevski, L., Castanon, G. & Tamil, L. S. *WDM fiber delay line buffer control for optical packet switching*. 2000. URL: <http://dx.doi.org/10.1117/12.401824>.
- [GK14] Gorja, P. & Kurapati, R. "Extending open vSwitch to L4-L7 service aware Open-Flow switch". *Advance Computing Conference (IACC), 2014 IEEE International*. 2014, pp. 343–347.
- [Har13] Harai, H. "Optical Packet and Circuit Integrated Networks". *Optical Network Design and Modeling (ONDM), 2013 17th International Conference on*. 2013, pp. 76–81.
- [HFFMW14] Harai, H., Furukawa, H., Fujikawa, K., Miyazawa, T. & Wada, N. "Optical Packet and Circuit Integrated Networks and Software Defined Networking Extension".

- J. Lightwave Technol.* **32.16** (2014), pp. 2751–2759. URL: <http://jlt.osa.org/abstract.cfm?URI=jlt-32-16-2751>.
- [Hoe12] Hoelzle, U. “OpenFlow at Google”. Open Networking Summit 2012, Santa Clara, California, 2012.
- [Hun07] Hunter, J. D. “Matplotlib: A 2D graphics environment”. *Computing in Science & Engineering* **9.3** (2007), pp. 90–95.
- [Jam10] James P.G. Sterbenz, Justin P. Rohrer, Egemen K. Cetinkaya, Mohammed J.F. Alenazi, Austin Cosner, and James Rolfe. *KU-Topview Network Topology Tool*. The University of Kansas. 2010.
- [Jue01] Jue, J. P. “Lightpath Establishment in Wavelength-Routed WDM Optical Networks”. *Optical Networks: Recent Advances*. Ed. by Ruan, L. & Du, D.-Z. Boston, MA: Springer US, 2001, pp. 99–122. URL: https://doi.org/10.1007/978-1-4613-0291-9_5.
- [KNFBR11] Knight, S., Nguyen, H. X., Falkner, N., Bowden, R. & Roughan, M. “The Internet Topology Zoo”. *IEEE Journal on Selected Areas in Communications* **29.9** (2011), pp. 1765–1775.
- [Net] *Linux Foundation - netem*. URL: http://www.linuxfoundation.org/collaborate/workgroups/networking/netem#Emulating_wide_area_network_delays (visited on 03/01/2014).
- [MABPPRST08] McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S. & Turner, J. “OpenFlow: enabling innovation in campus networks”. *SIGCOMM Comput. Commun. Rev.* **38.2** (2008), pp. 69–74. URL: <http://doi.acm.org/10.1145/1355734.1355746>.
- [MHUD16] Miyazawa, T., Harai, H., Udechukwu, R. & Dutta, R. “Experimental Demonstration of Interconnection Between ChoiceNet Economy Plane and Optical Packet and Circuit Integrated Network Control Plane”. *2016 IEEE High Performance Switching and Routing (Workshop Paper)*. 2016.
- [MHUD18] Miyazawa, T., Harai, H., Udechukwu, R. & Dutta, R. “Designing and experimentally demonstrating automatic OPS/OCS/OpenFlow network control driven by ChoiceNet provider”. *Photonic Network Communications* **36.1** (2018), pp. 43–55.
- [MMG08] Mu, M., Mauthe, A. & Garcia, F. “A Utility-Based QoS Model for Emerging Multimedia Applications”. *2008 The Second International Conference on Next Generation Mobile Applications, Services, and Technologies*. 2008, pp. 521–528.

- [Ope13] Open Networking Foundation. “OpenFlow Management and Configuration Protocol (OF-Config 1.1.1)”. 2013. URL: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config-1-1-1.pdf>.
- [QN15] Quinn, P. & Nadeau, T. *Problem Statement for Service Function Chaining*. RFC 7498. 2015. URL: <http://www.ietf.org/rfc/rfc7498.txt>.
- [Rou01] Rouskas, G. N. *Routing and Wavelength Assignment in Optical WDM Networks*. 2001.
- [RBCDGNW] Rouskas, G. N., Baldine, I., Calvert, K. L., Dutta, R., Griffioen, J., Nagurney, A. & Wolf, T. “ChoiceNet: Network Innovation Through Choice”. *Proceedings of ONDM 2013*.
- [Int] “Service-Aware Network Architecture Based on SDN, NFV, and Network Intelligence”. White paper. Intel, 2014. URL: <http://www.intel.com/content/www/us/en/communications/service-aware-network-architecture-sdn-nfv-paper.html>.
- [SGYACMP] Sherwood, R., Gibb, G., Yap, K., Appenzeller, G., Casado, M., McKeown, N. & Parulkar, G. “FlowVisor: A Network Virtualization Layer” ().
- [SNK12] Shin, M.-K., Nam, K.-H. & Kim, H.-J. “Software-defined networking (SDN): A reference architecture and open APIs”. *ICT Convergence (ICTC), 2012 International Conference on*. 2012, pp. 360–361.
- [SPYFGT13] Shin, S., Porras, P., Yegneswaran, V., Fong, M., Gu, G. & Tyson, M. “FRESCO: Modular Composable Security Services for Software-Defined Networks”. *Proceedings of the 20th Annual Network and Distributed System Security Symposium (NDSS’13)*. 2013.
- [Smo17] Smock, B. *k-shortest-paths*. 2017. URL: <https://github.com/bsmock/k-shortest-paths>.
- [SMJI13] Sun, L., Mkwawa, I.-H., Jammeh, E. & Ifeakor, E. *Guide to Voice and Video over IP: For Fixed and Mobile Networks*. Springer Publishing Company, Incorporated, 2013.
- [Tar71] Tarjan, R. “Depth-first search and linear graph algorithms”. *12th Annual Symposium on Switching and Automata Theory (swat 1971)*. 1971, pp. 114–121.
- [TASPCP14] Tomkos, I., Azodolmolky, S., Sole-Pareta, J., Careglio, D. & Palkopoulou, E. “A Tutorial on the Flexible Optical Networking Paradigm: State of the Art, Trends, and Research Challenges”. *Proceedings of the IEEE* **102.9** (2014), pp. 1317–1337.

- [UNOALS12] Ucar, I., Navarro-Ortiz, J., Ameigeiras, P. & Lopez-Soler, J. "Video Tester - A multiple-metric framework for video quality assessment over IP networks". *Broadband Multimedia Systems and Broadcasting (BMSB), 2012 IEEE International Symposium on*. 2012, pp. 1–5.
- [UBDR16] Udechukwu, R., Bhat, S., Dutta, R. & Rouskas, G. "Language of Choice: On Embedding Choice-related Semantics in a Realizable Protocol". *2016 IEEE 37th Sarnoff Symposium*. 2016, pp. 31–36.
- [UD14] Udechukwu, R. & Dutta, R. "Extending Openflow for Service Insertion and Payload Inspection". *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*. 2014, pp. 589–595.
- [UD15] Udechukwu, R. & Dutta, R. "Service definition semantics for optical services on a choice-based network". *Optical Network Design and Modeling (ONDM), 2015 International Conference on*. 2015, pp. 98–103.
- [UDMH17a] Udechukwu, R., Dutta, R., Miyazawa, T. & Harai, H. "Evaluating Different Pricing Algorithms for a Flexible Optical Choice-based Network". *2017 Proceedings of 26th International Conference on Computer Communications and Networks (ICCCN)*. 2017.
- [UDMH17b] Udechukwu, R., Dutta, R., Miyazawa, T. & Harai, H. "Toward Realizing Choice-based Co-Optimizable Networking Paradigm". *2017 IEEE 38th Sarnoff Symposium*. 2017.
- [UDMH18a] Udechukwu, R., Dutta, R., Miyazawa, T. & Harai, H. "Role of Provider Strategies and Granularity in a Circuit-Packet Optical Choice-based Network". *2018 Photonics in Switching and Computing (PSC)*. 2018, pp. 1–3.
- [UDMH18b] Udechukwu, R., Dutta, R., Miyazawa, T. & Harai, H. "Service Provisioning within a Choice-Based Network". *2018 International Conference on Recent Innovations in Electrical, Electronics & Communication Engineering (ICRIEECE)*. IEEE. 2018, pp. 2343–2346.
- [VLZJ11] Vahdat, A., Liu, H., Zhao, X. & Johnson, C. "The emerging optical data center". *2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference*. 2011, pp. 1–3.
- [VKC10] Veeraraghavan, M., Karol, M. & Clapp, G. "Optical Dynamic Circuit Services". *IEEE Communications Magazine* **48**.11 (2010), pp. 109–117.
- [VM14] Veeraraghavan, M. & Monga, I. "Broadening the scope of optical circuit networks". *Optical Network Design and Modeling, 2014 International Conference on*. 2014, pp. 275–280.

- [WGCDRBN14] Wolf, T., Griffioen, J., Calvert, K. L., Dutta, R., Rouskas, G. N., Baldin, I. & Nagurney, A. “ChoiceNet: Toward an Economy Plane for the Internet”. *SIGCOMM Comput. Commun. Rev.* **44.3** (2014), pp. 58–65. URL: <http://doi.acm.org/10.1145/2656877.2656886>.
- [WGCDRBN12] Wolf, T., Griffioen, J., Calvert, K. L., Dutta, R., Rouskas, G. N., Baldine, I. & Nagurney, A. “Choice as a principle in network architecture”. *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. SIGCOMM '12. Helsinki, Finland: ACM, 2012, pp. 105–106. URL: <http://doi.acm.org/10.1145/2342356.2342382>.
- [Yen71] Yen, J. Y. “Finding the K Shortest Loopless Paths in a Network”. *Management Science* **17.11** (1971), pp. 712–716. URL: <http://www.jstor.org/stable/2629312>.
- [Res] *YUV CIF reference videos (lossless H.264 encoded)*. URL: <http://www2.tkn.tu-berlin.de/research/evalvid/cif.html> (visited on 03/20/2014).