

Service-Oriented Multigranular Optical Network Architecture for Clouds

Georgios S. Zervas, Valerio Martini, Yixuan Qin, Eduard Escalona, Reza Nejabati, Dimitra Simeonidou, Fabio Baroncelli, Barbara Martini, Karim Torkmen, and Piero Castoldi

Abstract—This paper presents a novel service-oriented network architecture to bridge the informational gap between user applications and optical networks providing technology-agnostic multigranular optical network services for clouds. A mediation layer (service plane) between user applications and network control is proposed to facilitate a mapping process between user application requests and the network services. At the network level, a multigranular optical network (MGON) is proposed and implemented to support dynamic wavelength and subwavelength granularities with different transport formats [optical burst switched (OBS), optical burst transport (OBT)], reservation protocols (one-way, two-way), and different quality-of-service (QoS) levels per service type. The service-oriented multigranular optical network has been designed, implemented, and demonstrated on an experimental testbed. The testbed consists of service and network resource provisioning, service abstraction, and network resource virtualization. The service-to-network interoperation is provided by means of a gateway that maps service requests to technology-specific parameters and a common signaling channel for both service and network resource provisioning.

Index Terms—Optical networks; Service-oriented architecture; Cloud computing.

I. INTRODUCTION

Cloud computing is a service approach that provides IT services such as storage or computational capabilities enabled by distributed resources, such as data centers and clusters, to applications while hiding the relevant technology and implementation-specific details [1]. Cloud computing is different from grid computing, since the latter is oriented to distributed computation and it requires grid-specific applications.

Research on clouds has mainly focused on the IT infrastructure and technologies for both research and commercial purposes. Companies conducting clouds are deploying different cloud platforms (Amazon EC2 [1], Google App Engine [2], Microsoft Live Mesh [3], Sun Grid [4]), giving spe-

cial attention to factors such as CPU, memory, and storage provisioning. However, the success of cloud computing and the overall performance of the applications developed on it depends also on the physical network infrastructure that interconnects the distributed resources. In fact, clouds need proper consideration for network dynamics, such as capacity, quality of service (QoS), reservation, and complete integration and federation of the clouds. In such an environment, the transport infrastructure could be virtualized and treated as a type and part of the service [5,6]. The optical network infrastructure that follows a traditional architectural approach to creating and delivering network services may not be able to be part of the cloud's architecture and thus meet its requirements in a seamless way. Hence, it is vital to understand and redefine the role of optical networks, to not just to carry information created by the cloud, but also to be an active and integral element of it.

The principle of multigranular optical cross-connects (MG-OXCs) that can switch traffic at fiber, waveband, and wavelength granularities [7–9] has been proposed to reduce the cost and complexity of traditional OXCs. The MG-OXC is a key element for routing high-speed WDM data traffic in a multigranular optical network. The authors have complemented these studies by introducing a multigranular optical network (MGON) [10–12] that is able to support dynamic wavelength and subwavelength granularities with different QoS levels. By combining the two extensions, the full range of bandwidth granularities can be supported by a single OXC design. Moreover, it allows fast reconfigurability and flexibility on the electronic control of switching technologies while offering good cost–performance balance.

A MGON can support both optical circuit switched (OCS) and optical packet/burst switched (OPS/OBS) technologies with different transport formats and reservation protocols. In OCS networks, bandwidth granularity is at the wavelength level since one or more wavelengths are allocated to a connection, while connectivity between the source and destination is established using a two-way reservation, which is in general a time-consuming procedure. In OPS networks, one or more IP packets with similar attributes are aggregated in an optical packet and tagged with an optical header. This scheme does not need advance reservation, and it can provide on-demand one-way connectivity by imitating best-effort IP packet routing in the optical domain. OBS combines the advantages of OCS and OPS [13,14]. The fundamental premise is the separation of the control plane and data plane and the segregation of functionality within the appropriate domain (electronic or optical) [9].

Manuscript received August 14, 2009; revised July 2, 2010; accepted August 20, 2010; published September 30, 2010 (Doc. ID 115339).

G. S. Zervas (e-mail: gzerva@essex.ac.uk), Y. Qin, E. Escalona, R. Nejabati, and D. Simeonidou are with the High Performance Networks Group, School of Computer Science and Electronic Engineering, University of Essex, Colchester, CO4 3SQ, UK.

V. Martini, K. Torkmen, and P. Castoldi are with the CEIICP—Scuola Superiore Sant'Anna, Via G. Moruzzi 1, 56124 Pisa, Italy.

F. Baroncelli and B. Martini are with CNIT—National Laboratory of Photonic Networks, Via G. Moruzzi 1, 56124 Pisa, Italy.

Digital Object Identifier 10.1364/JOCN.2.000883

A cloud may benefit from a dynamic MGON able to provide on-demand and fine-grained connectivity to applications with a suitable level of network resource virtualization. The service-oriented optical network (SOON) framework [14] has been proposed to bridge the informational gap between applications and optical networks by introducing a mediation layer that facilitates a mapping process between user application requests and the network services. SOON aims at disjoining the parameters perceived by an end user from the technology-specific directives needed by network devices. This approach exists on clouds but only for the abstraction of IT infrastructure. SOON also enables automatic network configuration for establishing on-demand connection with different classes of service. As such, the SOON framework can be used as part of the infrastructure layer of the cloud computing stack [15] in order to oversee and deliver the network infrastructure as a service. The SOON framework can then enhance the agility of the cloud to re-provision technological infrastructures not only in the IT environment but also on the network too.

To address these challenges this work proposes a MGON able to deliver different levels of granularity as a service to clouds using the SOON framework. The SOON-enabled MGON allows dynamic interaction between application requests (network service) and network module configuration (edge node) by mapping application requests expressed in terms of QoS parameters (e.g., latency, bandwidth) to technology-specific attributes (e.g., burst size, wavelength, path, resource reservation method). A list of the benefits of proposed interoperable architectural and implementation solutions across the service plane, control plane, and data plane are shown below:

- Service-Oriented Optical Network (SOON)
 - Bridge the informational gap between the user application and the optical network
 - Directly offer resources (network and data) as service
 - Provide a mediation layer that facilitates a mapping process between user application requests and the network services
 - Disjoin the parameters perceived by an end user from the technology-specific directives needed by network devices
 - To facilitate intelligent discovery, reservation, and co-allocation of distributed resources across the network
- MGON
 - Bandwidth provisioning and switching capability at wavelength and subwavelength granularities
 - Agility and scalability of switching granularities providing a dynamic network infrastructure solution
 - Fast reconfigurability and flexibility on the electronic control of switching technologies
 - Cost-performance efficiency by offering an optimal balance between slow and fast switch fabric technologies
- SOON-enabled MGON advantages
 - Dynamic interaction between user request (network service) and network module configuration (edge node)
 - Network elements are able to recognize and process

application demands/requests

- Mapping of application requests to technology-specific attributes (burst size, latency, wavelength, path, resource reservation method)
- Service-aware connection establishment and QoS provisioning
- Fast provisioning and recovery

The remainder of this paper is organized as follows. First, the service-oriented multigranular network architecture is described and analyzed in Section II. This section describes the proposed SOON-MGON architecture and provides detailed information regarding the SOON internal signaling, the SOON-MGON integration, and the control and data plane packet format required for such integration. In Section III, the feasibility of MGON using different physical optical switches [microelectromechanical systems (MEMS) and acousto-optic] is demonstrated on an experimental test-bed. In addition, we report on the functionalities of SOON-MGON able to provide different network services for the QuadHD video-on-demand (VoD) application and also show its performance in terms of latency and jitter. Finally, Section IV concludes the paper.

II. SERVICE-ORIENTED MULTIGRANULAR NETWORK ARCHITECTURE

A. SOON-Enabled Multigranular Network Architecture

The SOON architecture, depicted in Fig. 1 introduces a technology-independent mediation layer, called the service plane (SP), that acts as a middleware between the application-layer and the network-layer. In particular, the SP composes and orchestrates the connectivity service pro-

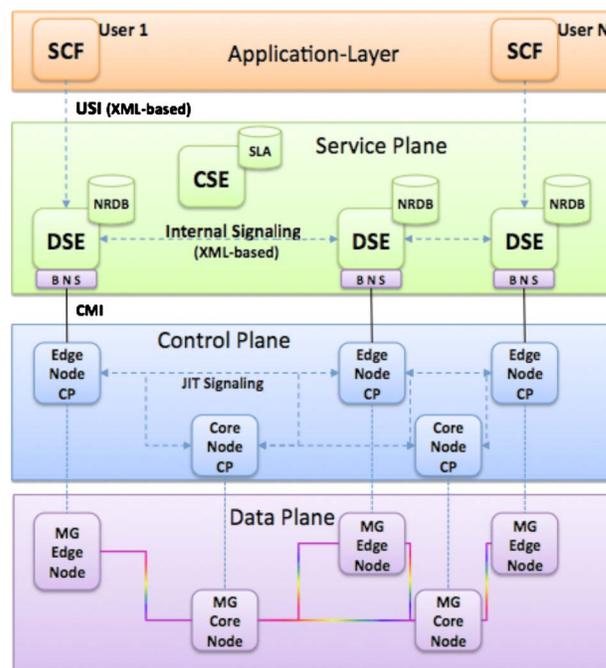


Fig. 1. (Color online) SOON architecture functional blocks.

vided by the control plane (CP) at the boundary of the transport network, thus providing technology-independent network services with a level of abstraction suitable for being invoked by an application. It also decouples the network technologies from existing and emerging network services and allows the CP to focus on connectivity provisioning services.

The SP [14] architecture is based on functional entities, called the centralized service element (CSE) and distributed service elements (DSEs). The CSE supports the SP functionalities regarding the identity of a customer application and the relevant contractual obligation with the network service provider. The SP does not require intervention of the network management system for service provisioning (i.e., on-request fashion), allowing the direct invocation of network services to the network (i.e., on-demand fashion). In this work, the application that is able to request connectivity services from the SP is represented by the service control function (SCF). A SCF interacts with the SP via the XML-based user-to-service interface (USI).

The SOON also has a compulsory technology-dependent sublayer called the broker network service (BNS) that interacts with specific devices and maps the SP command into specific technology directives. The BNS module works on one side with the SP entities and on the other side with the network via the control plane management interface (CMI). The CMI structure depends on the type of network transport service supported by SP operations. Usually, the CMI implementation is based on configuration protocols such as NETCONF [14]. In this work, the CMI is a proprietary interface suitable for the communication between the DSEs and the OBS devices. The SOON approach is as generic as possible and can be applied to other types of transport network systems by only changing the BNS module.

B. SOON Internal Signaling

To realize the above-mentioned functionalities, different signaling types and interfaces have been implemented to perform different functions. An automatic provisioning system has been implemented using two types of signaling among the DSEs called service provisioning signaling and background signaling.

1) Service Provisioning Signaling: The service provisioning signaling configures the network devices upon a network service-specific request (i.e., different message set per network service type) issued by an SCF. The service provisioning signaling may be subdivided into three phases: the *client request phase*, the *service correlation phase*, and the *network configuration phase*. Each phase involves a specific address space and QoS parameters with different granularity.

Figure 2 summarizes the parameter mapping from the SCF request through the SP and the set of network directives down to the MG edge node.

- a) **Client request phase:** This phase starts when an SCF sends a network-services client request message to a DSE to establish network service among a set of client networks. This message contains the following set of parameters:
 - SCF *identity parameters* are used to check the SCF identity and allow the establishment of a session be-

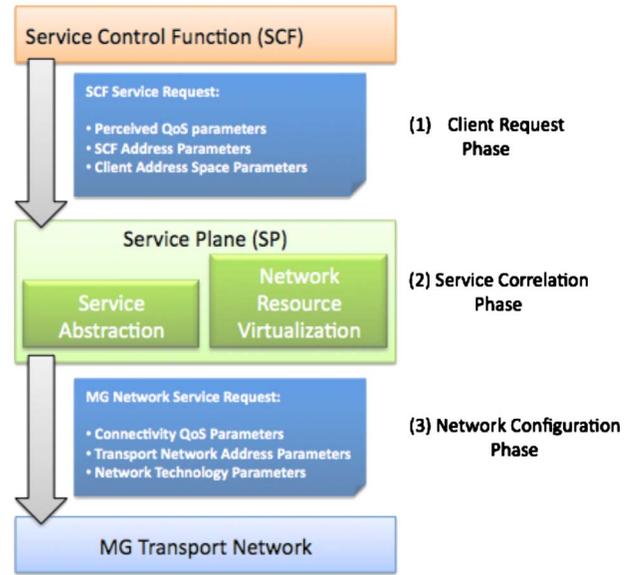


Fig. 2. (Color online) Three phases of provisioning signaling.

tween the SCF and the designed DSEs (e.g., the SCF-IP address, the SCF-TCP port used by send messages and receive response messages).

- *Perceived QoS parameters* are the QoS attributes that a customer is able to perceive and to monitor, using software or hardware probes, such as the bandwidth, the packet delay, or the packet jitter.
 - *Client address space parameters* represent the client address space and the CEs (customer edge nodes) involved in a network service and viewed by the provider.
- b) **Service correlation phase:** The DSE discovers the set of provider edges (PEs) involved in the network service provisioning from the list of CEs contained in the message provided during the client request phase. During the service correlation phase, service abstraction and network resource virtualization are performed.
 - The *service abstraction* maps a set of parameters that an application-layer module specifies for a particular service request into a set of specific parameters used by the network-layer for the identification and provisioning of that service.
 - The *network resource virtualization* aims at hiding the network technology details from the applications. The resource virtualization allows the application to request new services without dealing with network technology details. At the same time, virtualization improves network scalability and allows heterogeneous technologies to be used without the need to change the interaction with its customer applications.
 - c) **Network configuration phase:** The DSE communicates the parameters received during the service correlation phase to the DSEs, under the control of which the PE falls. In turn, each DSE uses these parameters to issue a set of node-specific network configuration request messages to the actual network element via the CMI. In this work the CMI is a proprietary protocol

transported by TCP messages.

2) *Background Signaling*: The background signaling is service independent and allows collecting and abstracting of the network status information (e.g., network technology and topology information) needed by the service provisioning signaling to perform dynamic correlation of the network service parameters among the network devices.

The background signaling is used to update the network resource database (NR-DB) within the DSEs in order to perform the service correlation phase previously described. The background signaling operation is periodically executed at regular intervals (typically 10 minutes or after the setup of a new network service).

As shown in Fig. 3, the background signaling consists of three phases:

- a) **Network resource discovery phase** (arrows 1, 2): Each DSE gathers at regular intervals the information regarding the identity of the PE and CEs attached to the controlled PE and the information related to the network service already established.
- b) **Network resource abstraction phase**: Each DSE maps the technology-specific information obtained during the resource discovery phase into the technology-independent information used by the DSEs to perform the service abstraction and network resource virtualization tasks. The information is stored in the NR-DB.
- c) **Information distribution phase** (arrows 3, 4): Each DSE distributes the information contained in its NR-DB to the CSE module that acts as a communication switch maintaining all the DSEs updated with the knowledge of the overall network information at the boundary of the optical burst switching and transport (OBST) nodes network.

C. SOON-MGON Interoperation for On-Demand Network Service Provisioning

To enable on-demand service provisioning within the SOON framework, some configuration operations must be

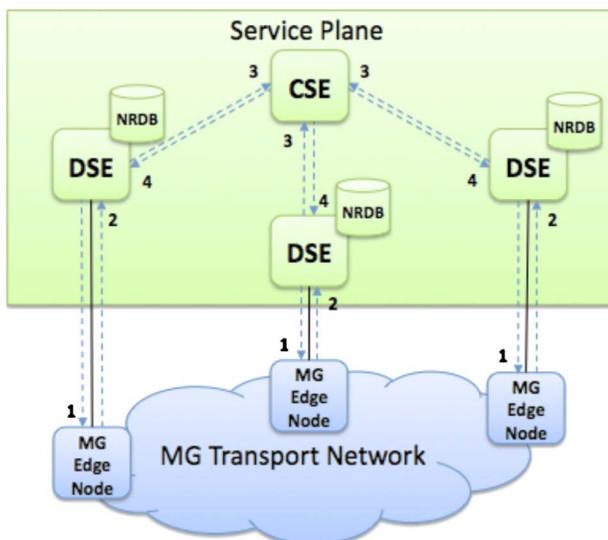


Fig. 3. (Color online) Background signaling.

performed between the different involved processes. The physical OBS information needed to populate the NR-BD is gathered via the CMI interface at regular intervals. This configuration allows the discovery of the physical interface used as well as the possible paths. Each DSE then associates that information with its address (Service_address) and the address of the relevant DSE.

The service-oriented multigranular network consists of edge and core nodes. The architecture of the edge node is displayed in Fig. 4. The main element of the architecture is the FPGA prototype (Virtex II-Pro), which is programmed to parse SOON messages and also perform a number of functionalities that will be discussed below.

As shown in Fig. 4, the DSE has 1GE connectivity to the FPGA board. Initially, the DSE sends a network configuration request (as part of phase 3 of service provisioning signaling) to the FPGA board via the CMI. The FPGA is then using a packet parser to identify and interpret all messages—SOON signaling messages and data packets. The SOON message called the “network services client request” is encapsulated within a burst control header (BCH) packet and forwarded over the MGON control plane to the destination DSE attached to a remote edge node.

The network configuration request message (SOON message) is used to trigger both the configuration of the edge node and the reservation of the end-to-end multigranular network service. This message is parsed and passed to the “burst scheduler parameter mapping” module. This module is responsible for mapping the SOON configuration request to different parameters on multigranular scheduling.

The SOON message is encapsulated on a TCP header using specific ports to provide technology-specific information. The integration of the SOON framework with MGON is carried out using a matrix that maps some type of services into a pool of parameters suitable from the edge node. Thus, the DSE module is in charge of the mapping process, and the FPGA interprets the matrix number received with the TCP encapsulated SOON message and configures the burst size, aggregation time (delay time), wavelength/path, and transport service used as a function of the TCP source port number received from the DSE (see Fig. 4). Figure 5 shows a timeline of the entire signaling flow from the service plane down to the OBS network burst signaling (BCH) up to the actual burst transmission.

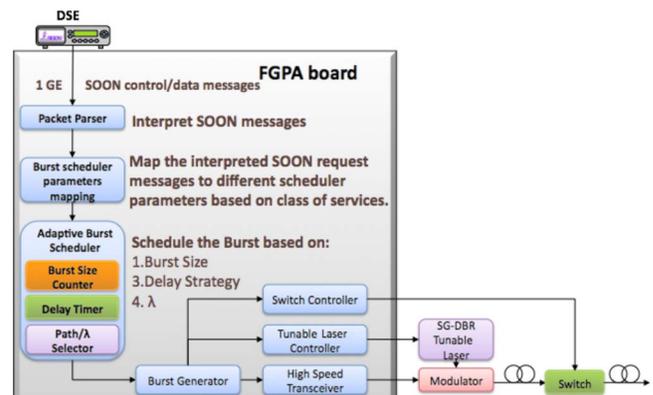


Fig. 4. (Color online) Edge node architecture.

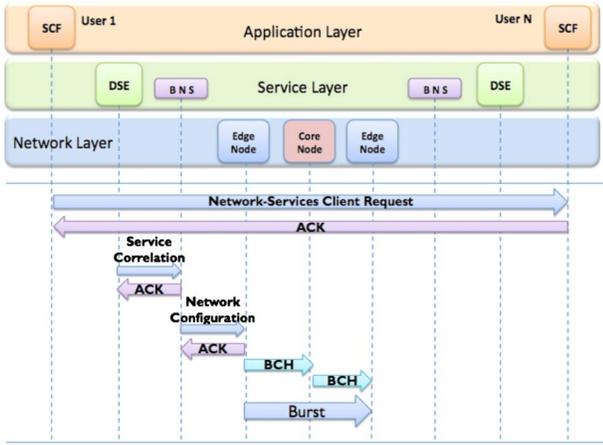


Fig. 5. (Color online) End-to-end signaling.

Table I provides detailed information on how client services are mapped to technology-specific parameters. Four different services are considered in this experiment: 1) QuadHD—VoD, 2) video-over-IP (VoIP), 3) video streaming, and 4) raw data transfer. Each of these has been mapped to different functions provided by the OBST testbed. The combination of the OBST functionalities depends on the application’s QoS requirements. Parameters such as the offered traffic, network status, and availability of network resources can also be included in the selection process. The burst size is affected by the limited internal RAM of the FPGA. The wavelength/path selection determines the number of hops and type of network resources used, which will affect the end-to-end delay and burst loss probability. Nevertheless, the testbed consists of two edge nodes, and thus there is no cause for contention at the core network and as a result no burst loss.

Two transport services are available for selection of the optical burst switching and the optical burst transport. The first one is based on the traditional OBS using one-way just-in-time (JIT) signaling per burst basis, whereas the second approach reserves a circuit for the duration of the service, allowing bursts to be transported over it.

D. Control and Data Plane Packet Formats Toward Service Orientation

To implement the functional modules of the SOON-enabled OBST network [16,17] and in particular the OBST edge router, new control and data plane packet/burst formats are proposed and defined as shown in Fig. 6. Regard-

TABLE I
CLIENT SERVICE MAPPED ON TECHNOLOGY-SPECIFIC PARAMETERS

Client Service	Technology-specific parameters			
Application	Burst Size (bytes)	Delay time (μs)	Wavelength/ Path	Transport Service
QHD - VoD	5120	5120	λ_2 / Path II	OBS
Voice	4096	4096	λ_1 / Path I	OBT
Video Streaming	3072	3072	λ_3 / Path III	OBT
Best Effort data	2048	2048	λ_4 / Path IV	OBT

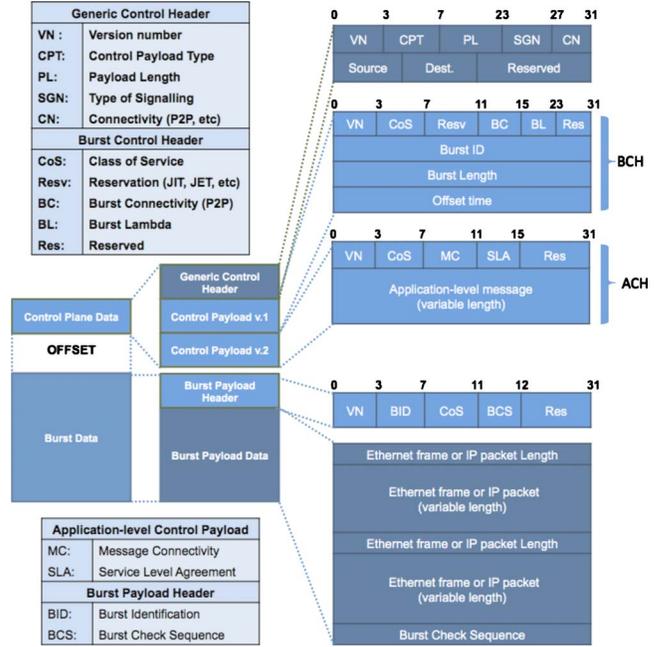


Fig. 6. (Color online) Control and data plane packet/burst formats.

ing the control plane, signaling information is divided into two main categories, the generic control header (GCH) and the control payload.

The **generic control header** is constructed to encapsulate any type of control payload required for different OBST network approaches (e.g., SOON-enabled OBS). It consists of 64 bits with the following fields:

- VN: *version number* required to identify the format of the control payload
- CPT: *control payload type* used to describe the encapsulated control payload
- PL: *payload length*
- SGN: the *type of signaling* used to represent in-band or out-of-band
- CN: the *connectivity* that indicates whether the control signal is sent to one or multiple destinations (point-to-point, multicast, anycast)
- Source
- Destination

The control payload format depends on the capabilities and enhancement of the OBST control plane. Here, two types are presented, namely, the BCH and the application control header (ACH).

The BCH is used on traditional OBST networks and is sent prior to the burst data in order to reserve network resources and establish end-to-end OBST connectivity. It consists of 128 bits with the following fields:

- VN: *version number*.
- CoS: *The class of service* is actually generated from the edge router classifier and is used to maintain the required CoS throughout the OBS network.
- Resv: *The reservation* block identifies the type of signaling mechanism [e.g., JIT, just-enough-time (JET), two-way] utilized to reserve the network resources.
- BL: *Burst lambda* identifies the lambda that the burst is carried on.

- The burst length and offset time are then required at each OBS core node in order to reserve and configure the switch for the appropriate time period.

The other type of control payload is the ACH, which is used to carry upper-layer messages and enrich the optical burst control plane with service information. The ACH has variable size and consists of the following fields:

- VN: It defines the format of the payload.
- CoS: same functionality as in BCH.
- Message connectivity (MC) is equivalent to the BC parameter of BCH.
- SLA: Service level agreement identifies the interaction between the service plane and the control plane. It is also used to encapsulate service-related messages generated by SOON and can be used to propagate/piggyback them over an OBST optical network.

After the detailed analysis of the SOON-OBS integration, signaling protocols, and packet format definition, Section III describes the testbed architecture/configuration, the experiments conducted, and the results measured.

III. NETWORK ARCHITECTURE AND TESTBED SETUP

The proposed service-oriented network architecture is based on SOON elements and OBST technologies utilizing service-aware edge and core routers interconnected by JIT SOON signaling. The architectural standpoint of the testbed is shown in Fig. 7, while the practical setup is shown in Fig. 8.

The SOON translates applications' service requests expressed in terms of perceived QoS and resources to a technology-specific pool at the edge of the OBST network. Through this capability of decoupling network technologies from services, the CP is unburdened of service-oriented functionalities and it can focus on the provisioning of connectivity services.

The SOON supports service abstraction and resource virtualization capabilities that allow a set of application-specific parameters to be mapped into a set of parameters used by the network for the actual configuration of a service while avoiding the exposition to applications of the technology-related details of network resources, such as the burst dimension, assembly time, network transport service (OBS, OBT), and end-to-end lightpath.

The SOON framework also has the ability to coordinate

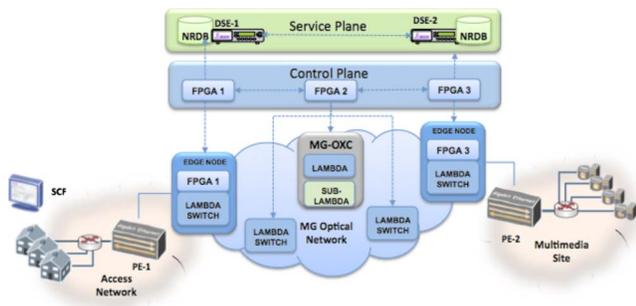


Fig. 7. (Color online) Architectural block diagram of the SOON-enabled OBST testbed.

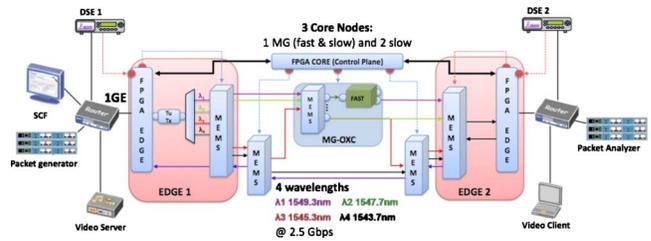


Fig. 8. (Color online) SOON-enabled OBST testbed setup.

the different OBS edge devices to create unidirectional and bidirectional end-to-end wavelength (OBT) and subwavelength (OBS) paths. This is obtained through an ad-hoc signaling protocol among DSEs, which can exchange the edge characteristic and reachability information. The DSE also manages the periodical update of information regarding the edge node devices used to solve the reachability of the source/destination through the network resource internal database (NR-DB). To enable the framework to interact with the OBS network a specific technology-dependent module installed in each DSE has been conceived, which translates the information received from the DSE into a set of specific directives comprehensible by the Ethernet-optical burst switching-transport (E-OBS-T) devices.

The service-aware edge OBS router, which utilizes a network processor and FPGA devices, is able to differentiate between service-layer messages (based on SOON), network requests, and data packets. In the case of SOON signaling messages, the edge router forwards them to the control plane. In the case of incoming SOON network requests, the burst aggregation scheduler is triggered to reflect service requirements into buffer size and the time thresholds. The SOON message is also used to decide on the network provisioning system of either OBS or OBT and then the appropriate lightpath (one out of a maximum of four). The OBS access control and provisioning system is created by aggregating bursts and in turn generating and transmitting a BCH ahead of time ($5 \mu s$) in order to configure the acousto-optic switch on a per-burst basis. OBT supports an end-to-end lightpath for the duration of the service by generating a BCH upon reception of a SOON message. The combination of lambda selection by controlling a SG-DBR tunable laser (at Edge 1) connected to a MEMS switch (for connectivity with all core nodes) together with the BCH processed at the core FPGA provides the end-to-end dynamic lightpath. Finally the data packets are buffered on aggregation first in, first out (FIFO) and then transmitted over different wavelength or subwavelength lightpaths. The OBST data plane transport mechanism is based on keep-alive messages in between burst transmission.

The service-aware core OBS router comprises three nodes all controlled by a centralized FPGA-based control plane module creating a meshed topology. The two nodes consist of MEMS switches (10 ms) and the third one of both a MEMS switch and an acousto-optic switch ($4 \mu s$) to form a multi-granular optical cross-connect (MG-OXC). The control plane module utilizes a network processor and an FPGA, and it can process and forward service-layer information on the fly as well as allocate switch resources for either OBS or OBT provisioning.

The OBS traffic is transported using a specific lightpath through the MG-OXC node and switched at the acousto-optic switch supporting service 1 (see Table I). The rest of the traffic is being transported over three different lightpaths having different numbers of hops based on OBST for the duration of the service.

Unlike the existing service-oriented architectures, which use the legacy IP network for carrying SOON signals and messages, in the proposed architecture SOON messages are carried in the optical domain with JIT signaling and within the GCH. The SCF application issues a request to the SOON framework for a specific network service (four network services). Then, the DSE element uses proprietary signaling to trigger the other involved DSE and configures the OBS devices. The SOON signaling is service specific and has different message sets for each type of provided service.

Figure 9 shows the overall time elapsed during the SOON network service provisioning process.

The first section (from the left) of the figure represents the SOON processing time of a service request before the start of the signaling (blue peak), while the second block represent the processing for the building of specific directives for the edge node configuration (pink peak) inferred from the user request. After the edge configuration the SOON gathers the ACK message from the DSE module and sends a service-provided ACK to the SCF service application (red line).

The SOON JIT control protocol performance has been evaluated by measuring the end-to-end service time, which includes the edge and core node parsing and forwarding time. This value is mostly dependent on the end host performance used for SOON elements and not the actual E-OBS-T testbed.

A. Experimental Results and Discussion

In this experiment, SOON service and connection establishment as well as high-definition video over OBST transmission are demonstrated. The SOON JIT messages encapsulated in GCHs are sent over the OBST control plane and the generated variable optical bursts are sent over an Ethernet-type data plane.

Figure 10 shows the delay and jitter of receiving 46 Mbps high-definition video over the E-OBS-T testbed, which has 200 Mbps background TCP traffic. Service 1 is defined and implemented in the edge node (FPGA) as best-effort class. Service 2 is defined and implemented in the edge node (FPGA) as timing critical class. In order to study the effect

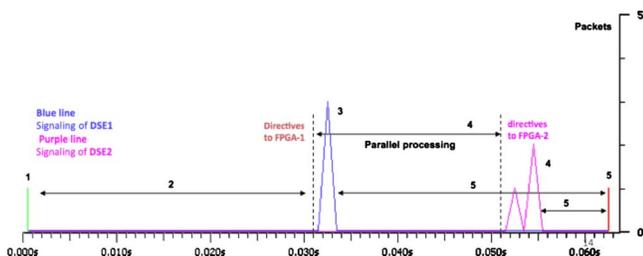


Fig. 9. (Color online) SOON edge-to-edge overall provisioning time.

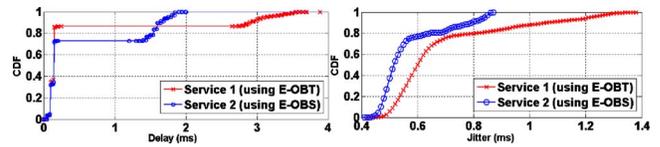


Fig. 10. (Color online) (a) Delay and (b) jitter for service 1 and 2.

of OBST on the real-time transmission of high-performance media, four prerecorded videos, with different qualities, were used in different streaming media scenarios across the OBST testbed network.

These videos varied from high definition with resolutions of 1280×720 and 1440×1080 and bit rates of 27 and 46 Mbps, respectively, to QuadHD of 2560×1600 and bit rates of 106 and 156 Mbps. TCP background traffic of around 200 Mbps was also generated through a traffic generator in order to emulate the current Internet traffic behavior (TCP and UDP data).

The aggregation developed is hybrid and combines both size and time threshold. Size and time thresholds are also dynamically changed per SOON service with a maximum size threshold of 5000 bytes and a time limit of 2 ms. Figure 10(a) shows that for service 1 (E-OBT) more than 95% of the UDP packets have a delay of less than 3 ms with a maximum delay of less than 4 ms, which is well within the acceptable level, and for service 2 the value is less than 1.8 ms. The horizontal lines that appear mean that no packets have delay less than a certain value due to the minimum burst aggregation time-out threshold. There are also several dumping dots close to zero. This is because when some packets enter the burst aggregation buffer, the aggregation timer is just about to time out. Figure 10(b) shows that for service 2 (OBS) the jitter also remains below 1.4 ms for 100% of the traffic and below 0.9 ms for service 1, also a well-accepted value.

Finally, Fig. 11 illustrates the physical layer performance (FER) of three different paths using different lambdas and different numbers and types of optical switches. Thus, B2B represents the back-to-back performance of the end-to-end testbed. Then, all different paths using different lambdas

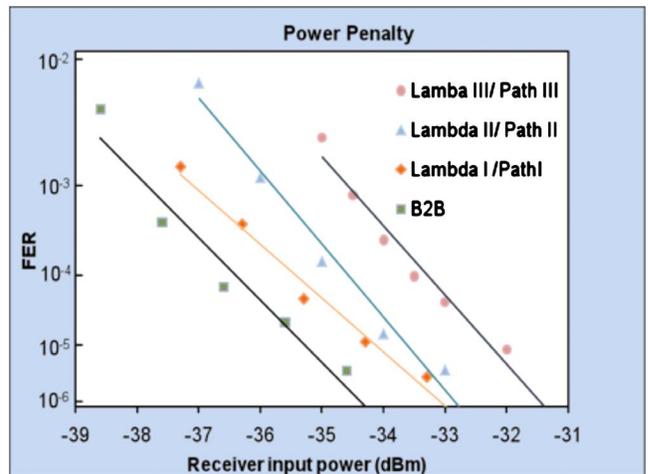


Fig. 11. (Color online) Frame error rate of all different network paths.

(e.g., Lambda I/Path I, Lambda II/Path II, Lambda III/Path III, all shown on both Figs. 8 and 11) experience a power penalty of 1.5–2.5 dB, which is well within acceptable levels considering the fact that the whole testbed is optimized once for all possible paths and the fact that each path is using a different number of MEMS switches (two up to five) or even a combination of MEMS switches and acousto-optic switches.

IV. CONCLUSIONS

A novel service-oriented multigranular optical network architecture has been presented as a candidate for clouds. Key architectural issues, protocols, packet formats, and techniques for a service-oriented MGON architecture were investigated. In particular, this work has explored a service plane that is able to facilitate a mapping process between user application requests and the network services through service abstraction and network resource virtualization. Then, the electronic gateway that has been developed in the edge node to translate user application requests to technology-dependent information (burst size, assembly time, wavelength/path, transport service) has been introduced. Also, a detailed description of the generic packet format required to carry service plane and control plane information over the same control channel (lambda) for cross-layer interoperation reasons has been provided. A testbed has been designed and implemented to support all above-mentioned functionalities, and experiments were conducted to show both the service plane performance (SOON provisioning time) in addition to application performance (QuadHD VoD) under different traffic conditions. Finally, the frame error rate (FER) has been experimentally measured for three different paths using different lambdas and different numbers and types of optical switches (MEMS and acousto-optic).

ACKNOWLEDGMENTS

The work described in this paper was carried out with the support of the BONE project (“Building the Future Optical Network in Europe”), a Network of Excellence funded by the European Commission through the 7th ICT-Framework Programme.

REFERENCES

- [1] Amazon Elastic Compute Cloud (EC2), <http://www.amazon.com/ec2/>, 2008.
- [2] Google App Engine, <http://appengine.google.com>, 2008.
- [3] Microsoft Live Mesh, <http://www.mesh.com>, 2008.
- [4] Sun network.com (Sun Grid), <http://www.network.com>, 2008.
- [5] C. E. Abosi, R. Nejabati, and D. Simeonidou, “A novel service composition mechanism for the future optical internet,” *J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. A106–A120, July 2009.
- [6] E. Grasa, A. Lopez, S. Figuerola, G. Junyent, and M. Savoie, “UCLPv2: a network virtualization framework built on web services,” *IEEE Commun. Mag.*, vol. 46, no. 3, pp. 126–134, Mar. 2008.
- [7] X. Cao, Y. Xiong, V. Anand, and C. Qiao, “Wavelength band switching in multi-granular all-optical networks,” *Proc. SPIE*, vol. 4874, pp. 198–210, July 2002.
- [8] P. H. Ho and H. T. Mouftah, “Routing and wavelength assign-

- ment with multi-granularity traffic in optical networks,” *J. Lightwave Technol.*, vol. 20, no. 8, pp. 1292–1303, Aug. 2002.
- [9] C. T. Politi, C. Matrakidis, A. Stavdas, D. Gavalas, and M. J. O’Mahony, “Single layer multigranular OXCs architecture with conversion capability and enhanced flexibility,” *J. Opt. Netw.*, vol. 5, no. 12, pp. 1002–1012, Dec. 2006.
- [10] G. S. Zervas, M. De Leenheer, L. Sadeghioon, D. Klondis, Y. Qin, R. Nejabati, D. Simeonidou, C. Develder, B. Dhoert, and P. Demeester, “Multi-granular optical cross-connect: design, analysis and demonstration,” *J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 69–84, June 2009.
- [11] R. Nejabati, G. Zervas, G. Zarris, Y. Qin, E. Escalona, M. O’Mahony, and D. Simeonidou, “A multi-granular optical router for future networks,” *J. Opt. Netw.*, vol. 7, no. 11, pp. 914–927, Nov. 2008.
- [12] M. Savi, G. Zervas, Y. Qin, V. Martini, C. Raffaelli, F. Baroncelli, B. Martini, P. Castoldi, R. Nejabati, and D. Simeonidou, “Data-plane architectures for multi-granular OBS network,” in *Optical Fiber Communication Conf.*, San Diego, 2009, paper OML5.
- [13] C. Qiao and M. Yoo, “Optical burst switching (OBS)—a new paradigm for an optical Internet,” *J. High Speed Netw.*, vol. 8, no. 1, pp. 69–84, 1999.
- [14] B. Martini, V. Martini, F. Baroncelli, K. Torkman, and P. Castoldi, “Application-driven control of resources in multiservice optical networks,” *J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. A270–A283, July 2009.
- [15] S. Johnston, “Taxonomy: the 6 layer cloud computing stack,” <http://samj.net/2008/09/taxonomy-6-layer-cloud-computing-stack.html>.
- [16] Y. Sun, T. Hashiguchi, V. Q. Minh, X. Wang, H. Morikawa, and T. Aoyama, “A burst-switched photonic network testbed: its architectures, protocols and experiments,” *IEICE Trans. Commun.*, vol. E88-B, no. 10, pp. 3864–3873, Oct. 2005.
- [17] S. Ovadia, C. Maciocco, M. Panicia, and R. Rajaduray, “Photonic burst switching (PBS) architecture for hop and span-constrained optical networks,” *IEEE Commun. Mag.*, vol. 41, no. 11, pp. S24–S32, Nov. 2003.



Georgios S. Zervas was awarded the M.Eng. degree in electronic and telecommunication systems engineering with distinction and the Ph.D. degree at the University of Essex (UK) in 2003 and 2008, respectively. He is a Research Fellow in the High Performance Networks Group at the University of Essex involved in the current and past EC-funded projects MAINS, STRONGEST, GEANT3, BONE, Phosphorus MUFINS, and e-Photon/One. He is

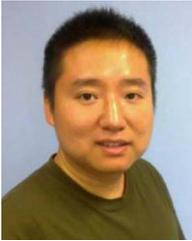
author and co-author of over 60 papers in international journals and conferences. His research interests include high-speed optoelectronic router design, subwavelength networks (e.g., OBS), GMPLS, cognitive networks, and grid/cloud networks. He is also involved in standardization activities in the Open Grid Forum (OGF).



Valerio Martini received his master’s degree (M.Sc.) in telecommunication engineering in March 2005 from Pisa University (Italy) defending his thesis on QoS optimization in virtual LAN (VLAN) protocol (IEEE 802.1Q) using Opnet Network simulator in the Network Laboratory of the Information Technology Department. He has been a Research Fellow Academy Visitor at the KTH, Royal Institute of Technology, in Stockholm (Sweden) in 2007 and at the University of Essex, Colchester, London (UK), in 2008. His research interests include network virtualization, L2/L3 VPN service technologies, service virtualization in (G)MPLS networks, and cloud computing. He received the Ph.D. degree in telecommunications

university of Essex, Colchester, London (UK), in 2008. His research interests include network virtualization, L2/L3 VPN service technologies, service virtualization in (G)MPLS networks, and cloud computing. He received the Ph.D. degree in telecommunications

from “Scuola Superiore Sant’Anna” in March 2009 working at the Integrated Research Center for Photonic Networks and Technology (CEIIC), Pisa (Italy). At the moment he holds the position of System Engineer at Juniper Networks Amsterdam.



Yixuan Qin (Ph.D. in electronic system engineering) is a Senior Research Officer in the High Performance Networks Group at the University of Essex involved in the EC-funded projects MUFINS, e-Photon/One+, BONE, DICONET, MAINS, and STRONGEST since 2007. His research interests include high-speed digital system design, embedded system design, hardware accelerated computing/networking, flexible networks, passive optical networks, optical

burst switching, impairment-aware-based routing, and GMPLS networks. He has over 20 publications including journals, conference papers, and a book chapter.



Eduard Escalona has been a Research Member within the High Performance Networks Group at the University of Essex since 2007. He obtained his M.Sc. degree in telecommunications engineering (2003) at the Universitat Politècnica de Catalunya (UPC). He has participated in several European projects such as IST LION, NOBEL, CARISMA, and PHOSPHORUS. He is currently involved in the GÉANT3 and FP7 GEYSERS projects. His main research areas

are management and control of optical networks, focused on control plane issues and service-oriented architectures.



Reza Nejabati is currently an Academic Member of the High Performance Networks Group in the School of Computer Science and Electronic Engineering, University of Essex. He has over ten years of academic and industrial experience in the field of telecommunication and computer science. The main current areas of his research interests are the application of ultra-high-speed network technologies for the future Internet, design and control issues for

service-oriented and application-aware networks, network virtualization as well as high-performance network architecture and technologies for e-science applications. He has more than 80 publications in these areas.



Dimitra Simeonidou is the Head of the High Performance Networks Group, which includes the Photonics Networks Laboratory and the newly established Network Media Laboratory, at the University of Essex, UK. She joined Essex in 1998 (previously with Alcatel Submarine Networks). While at Alcatel, she held the post of Senior Principle Engineer and contributed to the introduction of WDM in long-haul submarine links and pioneered the design and deployment of optical add-drop multiplexers. At Essex, she is leading a group of 30 researchers and Ph.D. students and she is involved in numerous national and international research projects. Her research is focusing in the fields of optical networks, grid and cloud computing, and the future Internet. She is author and coauthor of over 350 papers, 11 patents, and several standardization documents.



Fabio Baroncelli (M’05) received the master’s degree in physics in 1997 from the University of Pisa (Italy). He worked for Synapsis as a Web Software Engineer and for Marconi Communications as a Software Engineer in DWDM technology and network management. Since 2003 he has been a Researcher at the “Centro di Eccellenza per l’Ingegneria dell’Informazione, della Comunicazione e della Percezione” (CEIICP) in Pisa (Italy). His main research interests

include optical control and management plane design, cloud computing architectures, and hybrid MANET.



Barbara Martini received the master’s degree in electronic engineering in 1999 from the University of Florence (Italy). She joined Italtel from university working as a Hardware Engineer on network device driver design and TCP/IP stack protocols and then joined Marconi Communications in the summer of 2000 as a Software Engineer involved in network management software design in DWDM equipment. Since 2003 she has been a Research Engineer at

the CNIT National Laboratory of Photonics Networks located in Pisa (Italy). Her main research interests include network management system design, GMPLS optical control plane architectures, and service platform architectures in next-generation networks. Moreover, she also carries out researches in advanced security supports for multidomain networks.



Karim Torkmen was born in Tunis (Tunisia) in 1982. He received the National Diploma of Engineer in telecommunications from the High School of Communications of Tunis (Sup’Com), Tunis (Tunisia), in 2006. He achieved the degree of International Master on Communication Networks Engineering (IMCNE) at Scuola Superiore Sant’Anna, Pisa (Italy), in 2007. Currently he is a Ph.D. student at the Centre of Excellence for Information, Communication and

Perception Engineering (CEIICP) of Scuola Superiore Sant’Anna, Pisa (Italy). His main research interests include network service provisioning to multimedia applications and mobility support in IP networks.



Piero Castoldi (Ph.D. in information technology) has been an Associate Professor at Scuola Superiore Sant’Anna, Pisa, Italy, since 2001. He spent abroad at Princeton University (USA) overall about two years in 1996, 1997, 1999, and 2000, and in 2001 he visited for two months the University of Texas at Dallas, USA. He has also served as a Project Manager of many projects of the Interuniversity National Consortium for Telecommunications (CNIT), and since

January 2005 he has been the Director of the CNIT National Laboratory of Photonic Networks. He is also currently the Leader of the “Networks and Services” research area at the Center of Excellence for Networks Engineering at Scuola Superiore Sant’Anna. His research interests cover telecommunications networks and systems, both wired and wireless, and more recently reliability, switching paradigms, and control of optical networks, including application-network cooperation mechanisms, in particular for grids and clouds. He is author of more than 200 publications in international journals and conference proceedings.