

# Planning and Operating Flexible Optical Networks: Algorithmic Issues and Tools

Aristotelis Kretsis, Konstantinos Christodoulopoulos, Panagiotis Kokkinos, and Emmanouel Varvarigos,  
University of Patras

## ABSTRACT

We discuss the basic algorithmic issues in planning and operating flexible optical networks, highlighting the challenges and differences from fixed-grid WDM networks. We also overview commercially available planning and operation tools, which (expectedly) lack support for flexible optical networks at the moment, as flexible technologies are not yet deployed in real networks. To fill this need we present the Mantis tool. Mantis' modular architecture provides fast running times, efficient and fault-tolerant usage of computation resources, and execution as both a desktop application and a cloud software-as-a-service.

## INTRODUCTION

The continuous growth of consumers' IP traffic, fed by the generalization of broadband access (through digital subscriber line and fiber to the home) and the emerging rich-content high-rate and bursty applications, such as video on demand, HDTV, and cloud computing, can only be met with the abundant capacity provided by optical transport networks. For the future, it is expected that the traffic will not only increase in volume (traffic increase of 34 percent on average per year [1]) but will also exhibit high burstiness, resulting in large variations over time and direction.

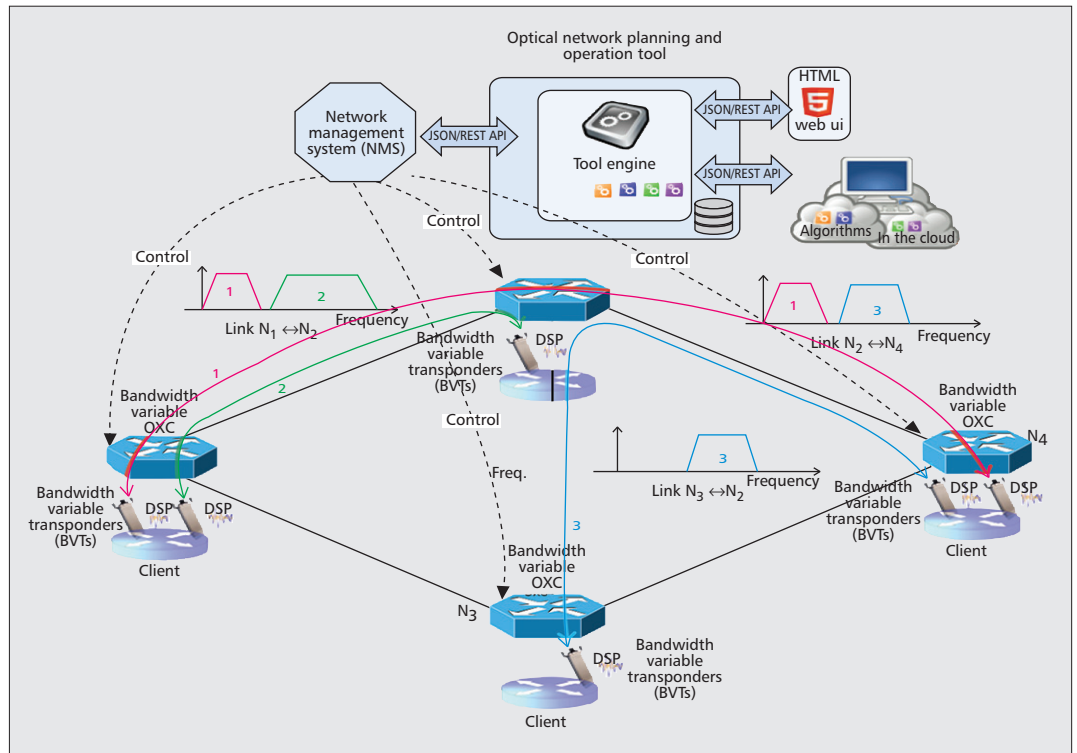
Recent research efforts on optical networks have focused on architectures that support variable spectrum connections as a way to increase spectral efficiency and reduce costs. Flexible (or elastic) optical networks appear as a promising technology for meeting the requirements of next generation networks that will span across both the core and metro segments, and potentially also across the access, all the way to the end user. A flexible network is based on the *flex-grid* technology, which migrates from the *fixed* 100 or 50 GHz grid that traditional (dense) wavelength-division multiplexing, (D)WDM, networks utilize [2]. Flex-grid has granularity of 12.5 GHz, standardized by the International Telecommunication Union (ITU, G.694.1), and can combine the spectrum units, referred to as *slots*, to create wider channels on an as needed basis. Flexible networks [3] are built using bandwidth variable optical switches that are configured to create optical paths of sufficient spectrum slots. We

refer to such a connection as a *flexpath*, a variation of the word *lightpath* used in standard WDM networks. Bandwidth variable switches operate in a *transparent* manner for transit (bypassing) traffic that is switched while remaining in the optical domain.

Flexible networks in addition to flex-grid switches assume the use of bandwidth variable transponders (BVTs) [3]. Various BVT implementations exist [4], employing single- or multi-carrier transmission schemes, and usually having some sort of digital signal processing (DSP) capabilities at the receiver but also at the transmitter side. Several transmission parameters can be controlled in a BVT, including the baud rate (symbols processed per second), the modulation format (number of bits encoded per symbol), the forward error correction (FEC) used, the spectrum slots employed, and the useful bit rate. Since transmission parameters are controllable, the term *software defined optics* has also recently been used, implying that optical networks, which currently rely on the slowly changing circuit switching paradigm, become more dynamic. Deciding the transmission parameters is quite complicated since physical layer impairments (PLIs) such as noise, dispersion, interference, and nonlinear effects accumulate and deteriorate the quality of transmission (QoT) of the flexpaths. In particular, the QoT of a flexpath depends on its BVT transmission parameters, the guardband used from its spectrum-adjacent flexpaths, and their transmission parameters.

Flexible networks built out of flex-grid switches and BVTs create flexpaths with as many spectrum slots as actually needed; thus, the huge but still limited optical bandwidth is allocated in an efficient manner. In addition, thanks to their adaptability, flexible networks enable the dynamic sharing of the network resources, thus constituting a fine-grained cost- and power-efficient network able to carry traffic that may vary in time, direction, and magnitude. Network equipment vendors and operators are already looking into this technology [3, 5]. A key enabler for the introduction of flexible technology will be the creation of network planning and operation tools that will control such networks, evaluate the benefits, and enable the adaptation of the candidate technology in real systems. Figure 1 presents an example of a flexible network that is

After the network is deployed we pass to the network operation phase. In this phase capacity demands are generally processed one or a set at a time, and the traffic is accommodated using the equipment already deployed in the network or when necessary some additional equipment.



**Figure 1.** A flexible optical network, consisting of bandwidth variable optical switches (BV-OXCs) and transponders (BVTs), that is controlled by an optical network planning and operation tool.

controlled by such a centralized network planning and operation tool.

Network planning is an important operation for all kinds of networks, and particularly for the optical networks constituting the core that interconnects all other networks together. The multiple degrees of freedom present in flexible networks, the adaptability of the BVT, and the interdependence between the transmission parameters and the physical layer make connection establishment more complicated than in fixed-grid WDM networks. Moreover, the adaptability of flexible networks brings the optical layer closer to the IP layer, enabling a cross-layer consideration of the two layers.

The emergence of flexible networks and new application methodologies (clouds, software as a service — SaaS, social by design) require not simply the extension of existing network planning tools, currently available from several major players, but the implementation of new ones. In this article, we describe the basic algorithmic challenges that a network planning and operational tool for flexible optical networks should address and present a related such tool, called Mantis, that contains a rich library of algorithms. Mantis was designed and implemented so as to accommodate both desktop and cloud execution.

## ALGORITHMIC ISSUES IN FLEXIBLE OPTICAL NETWORKS

We can broadly separate a network's life cycle in two phases. The planning phase typically occurs before a network is deployed, and focuses on how to better accommodate the current and foreseen

network traffic. Greenfield planning refers to the case where the network, including its physical topology, is designed from scratch, and can be very demanding, since many costs, such as the cost of digging ducts to lay the fiber cables, and the cost of the fibers, amplifiers, switches, transponders, and regenerators to be installed, have to be accounted for. More common in optical transport networks is to plan the network assuming that the fibers are already installed, which is the reality for tier 1 and tier 2 operators. In this case the topology is given, and the purpose is to decide the equipment (transponders, regenerators, and switches) to be purchased and where/how to deploy it. The objective is to minimize the equipment cost and the resources used for serving the given traffic (saving resources for future use or, in case of infrastructure leasing, for cost minimization).

After the network is deployed we pass to the network operation phase. In this phase capacity demands are generally processed one or a set at a time (e.g., a network operator can provision once per quarter or per year), and the traffic is accommodated using the equipment already deployed in the network or, when necessary, some additional equipment (e.g. transponders/regenerators) that is deployed. Therefore, the operation process must take into account any constraints posed by the current state of the network.

In the following we highlight the challenges and differences in planning and operating flexible and fixed-grid WDM networks from an algorithmic perspective, focusing on three topics:

- Accounting for physical layer impairments
- Routing, spectrum allocation, and choice of modulation format
- Dynamic network operation

## ACCOUNTING FOR PHYSICAL LAYER IMPAIRMENTS

Optical networks have evolved over recent years from opaque (point-to-point) to transparent networks, as a way to reduce capital and operating expenditures (CAPEX and OPEX). In the latter case, optical switches are configured to transparently handle transit traffic; the signal remains in the optical domain, saving on the cost of transponders and electronic processing equipment used in the past to terminate and retransmit traffic at intermediate hops. Since optical connections may span over many and long links, physical layer impairments (PLIs), such as noise, dispersion, interference, and nonlinear effects accumulate and affect the quality of transmission (QoT). Accounting for PLIs is a challenge for algorithm designers, especially when it comes to the interdependencies introduced between the physical and the other layers.

PLIs affect both fixed-grid WDM and flexible networks, but there are distinct differences between the two cases. Flexible networks are expected to use coherent detection and DSP, implying that certain impairments, particularly those related to dispersion, will be substantially reduced or fully compensated. However, the additional degrees of flexibility available in flexible networks make the minimization of these effects more complicated from an algorithmic perspective. On the other hand in WDM networks, PLIs, even though more significant, can be accounted for quite accurately, since fewer parameters are involved (non-tunable transponders and constant guardband) and analytical models can successfully capture these effects.

To formulate the PLI effects and the transponders' tunability in a flexible network we assume that each flexpath has a specific optical reach, defined as the distance it can transmit at with acceptable QoT (e.g., bit error rate, BER). The optical reach depends not only on the flexpath's transmission configuration, but also on the presence of adjacent interfering flexpaths, their transmission configurations and guardbands used. The number of combinations of possible configurations can be huge; also, PLI analytical models may not capture all effects, or experimental measurements may be limited for some of the options. So, it seems that the only viable solution is to resort to some sort of simplification that captures PLIs in a coarser but safe manner (e.g., add a safety margin or take worst-case scenarios), reducing the parameters and the solution space while trying not to eliminate good solutions.

In the algorithms implemented in the Mantis tool we used such a simplification that fits well with the above described requirements. The transmission reach of a connection for a given transmission configuration and guardband distance is calculated assuming worst case interference by its adjacent connections. To be more specific, we define a *physical feasibility function*  $l = f_c(r, b, g)$  that captures PLIs and can be obtained experimentally or through analytical models. This function gives the transmission distance  $l$  (in kilometers) at which the transponder of cost  $c$  can transmit with acceptable QoT when

transmitting at  $r$  Gb/s using bandwidth of  $b$  spectrum slots and a guardband of  $g$  spectrum slots from its adjacent spectrum connections. Note that defining the rate  $r$  and spectrum  $b$  incorporates the choice of the modulation format used.

Using the functions  $f_c$  of the available transponders we define (*reach-rate-spectrum-guardband-cost*) transmission tuples, corresponding to feasible configurations of the transponders. The term "feasible" is used to signify that the tuple definition incorporates PLI limitations, while the cost parameter is used when there are transponders of different capabilities and costs. The above definition is very general and can be used to describe any type of flexible or even fixed-grid optical network. For example, the mixed-line-rate (MLR) fixed-grid network of [8] that employs 10/40/100 Gb/s transponders can be represented with the following transmission tuples:

- (3000 km, 10 Gb/s, 50 GHz, 0 GHz, 1 cost unit)
- (1600 km, 40 Gb/s, 50 GHz, 0 GHz, 3 cost units)
- (800 km, 100 Gb/s, 50 GHz, 0 GHz, 6 cost units)

where the cost of the 10 Gb/s transponder is taken as 1 cost unit.

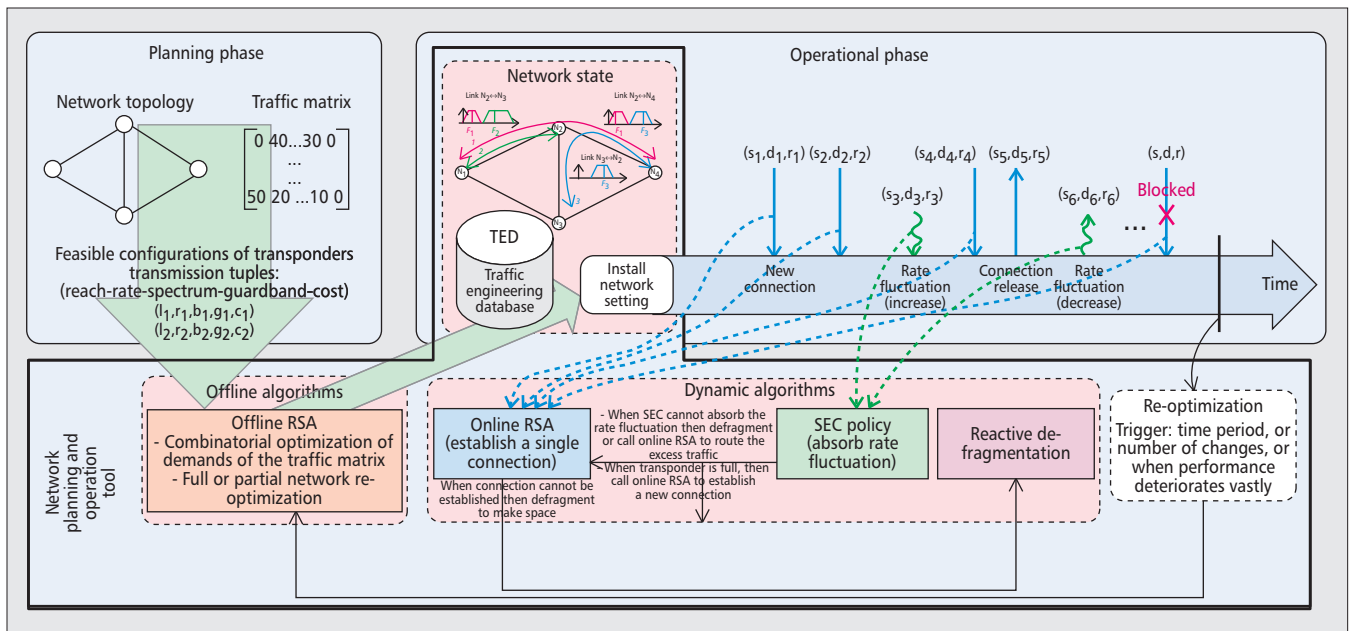
Using the above methodology we can enumerate the feasible transmission options of the BVTs so as to incorporate PLIs. The planning and operation algorithms of Mantis take these as input, and select the appropriate BVT configuration and allocate route and spectrum, in a way that satisfies QoT constraints by construction.

## ROUTING AND SPECTRUM ALLOCATION

The problem of establishing connections in fixed-grid WDM networks is typically referred to as the Routing and Wavelength Assignment (RWA) problem, known to be NP-complete. Connection (flexpath) establishment in flexible networks is more complicated for several reasons. First, in contrast to WDM networks where each connection is assigned a single wavelength, in flexible networks spectrum slots can be combined to form variable-width channels, leading to the so-called routing and spectrum allocation (RSA) problem. Additionally, the transponders' (BVTs) adaptability yields many transmission options, each with different transmission reach and spectrum used. Algorithms that try to capture, at varying degrees, the problem of jointly allocating recourses and selecting the transmission configurations are referred to as distance adaptive spectrum allocation or routing, modulation level, and spectrum allocation (RMLSA) algorithms.

Another difference between the RWA and RSA problems relates to the number of connections served by a transponder. Typically in WDM networks a transponder is allocated to a single connection. In flexible networks, since the BVTs are powerful, it is hard to fully utilize them for all connections. So *sliceable* BVTs (S-BVTs), which can be shared to serve more than one connection, are also considered as candidate solutions [3]. S-BVTs introduce additional challenges: how to assign multiple flexpaths to S-BVTs, keep them highly utilized, and reduce

*The problem of establishing connections in fixed-grid WDM networks is typically referred to as the Routing and Wavelength Assignment (RWA) problem, known to be NP-complete. Connection (flexpath) establishment in flexible networks is more complicated for several reasons.*



**Figure 2.** Flow diagram of a generic approach to operate a flexible network. The RSA algorithm performs routing and spectrum allocation for long- and medium-term provisioning, while SEC policy accommodates short-term traffic fluctuations.

their number and total cost. Thus, instead of point-to-point transmissions, algorithms that consider point-to-multipoint transmissions have to be developed, and ideas from multicast routing can be used.

Both RWA and RSA include as a subproblem the placement of regenerators in the network. Given the limited transmission reach due to PLIs, regenerators are used to establish lengthy connections. However, in contrast to WDM networks where the capabilities of the transponders and regenerators are given, in flexible networks the transmission reach depends on the transmission configuration of the tunable BVT and can be controlled. Thus, regenerator placement in flexible networks also depends on the choices made in choosing the BVT configurations, making it more complicated than in fixed-grid networks.

### DYNAMIC NETWORK OPERATION

The WDM network is typically initiated with an offline/planning algorithm assuming an oversubscribed traffic matrix, meaning that actual traffic is on average much lower than the traffic described in the matrix, to absorb short term fluctuations (e.g., daily cycles) and avoid frequent network upgrades. Of course, in the long term the traffic will increase and will eventually require network upgrades. Thus, WDM networks are operated in an incremental manner, with new connections added sporadically, when utilization between endpoints exceeds a certain percentage, while existing connections are rarely (if ever) terminated. This practice is rather different than the dynamic scenarios usually considered in literature algorithms for WDM networks that assume the dynamic establishment and release of connections and measure the blocking probability over time. It is actually very seldom that connections are torn down by tier 1 and tier

2 operators. Thus, a more appropriate model would be to incrementally add connections and observe when the first connection blocking occurs. Flexible networks using adjustable transponders (BVTs) require a different approach as their operation could be more dynamic, bringing the optical layer closer to the IP layer, for example, by enabling the IP layer to “dial” the bandwidth that it uses.

We envision a flexible network where dynamic traffic variation is accommodated at two different levels. The first level is the establishment of new connections, as in fixed-grid networks. Given the high rate that flexible BVTs are expected to transmit (designs of 400 Gb/s or 1 Tb/s have appeared in the literature [4]), relatively long periods of time will pass until a new connection (traffic increase) requiring a new transponder is encountered, probably longer than in WDM systems due to the higher rate of BVTs. A second level is to absorb rate changes by adapting the BVT, e.g., tuning the modulation format and/or the number of spectrum slots, a feature not available in WDM systems.

Figure 2 describes a generic approach for operating a flexible network. The offline RSA algorithm used to initialize the network, or the dynamic RSA algorithm that subsequently adds flexpaths, assigns to each flexpath a path and a reference frequency. A flexpath occupies a certain amount of spectrum slots around that reference frequency, and traffic variations can be absorbed by the BVTs by tuning the modulation format or expanding/contracting the spectrum they use. Slots that are freed by a flexpath can be assigned to different flexpaths at different time instants, obtaining statistical multiplexing gains. To enable the dynamic sharing of spectrum, we need Spectrum Expansion/Contraction (SEC) policies [10] to regulate the way this is performed. An example of such a policy would

Tools	IP/MPLS planning	DWDM planning	Flex-grid planning
Cisco - Cariden www.cariden.com/products/mate-product-family/	√		
Opnet - NetOne network planner www.opnet.com/solutions/network_management/spguru_network/planner/	√		
Opnet - NetOne transport planner www.opnet.com/solutions/network_management/transport_planner.html		√	→
Wandl - IP/MPLS www.wandl.com/html/mplsview/ipmplsview.php	√		
Wandl - NPAT www.wandl.com/html/npat.php		√	→
Aria - IP/MPLS - TE operational planning www.aria-networks.com/solutions/ip-mpls-te-operational-planning	√		
VPI - Multi-layer transport optimization www.vpisystems.com/solutions/multi-layer-transport-planning---optimization		√	→
Infinera - network planning system www.infinera.com/products/mgmtsuite/html	√		
Nokia Siemens networks SURPASS www.nokiasiemensnetworks.com/portfolio/products/transport-solutions/network-management-planning		√	→
Cisco - Transport planner www.cisco.com/en/US/prod/collateral/optical/ps5726/ps11348/data_sheet_c78-658849.html		√	→

Figure 3. Network planning tools overview.

A proactive approach can be an RSA algorithm that avoids fragmentation by trying to leave spectrum voids that are more easily usable, while a reactive approach can be a special defragmentation algorithm that re-optimizes the network and frees spectrum for new connections.

be to increase and decrease the spectrum slots in a symmetric way, i.e., alternate between expanding towards the higher and lower spectrum slots of the reference frequency of the flexpath, until it reaches a slot already occupied by another connection, at which point we allocate slots only from the other side. An RSA operation is performed again when a SEC policy cannot absorb traffic variations by granting additional spectrum slots, or when the requested rate exceeds the transponder capabilities. Then, RSA algorithm is called to route the excess traffic over a different flexpath, or reroute the entire flexpath (to save in guardbands).

After establishing and tearing down multiple flexpaths the spectrum slowly becomes fragmented, reducing its ability to accommodate new connections [11]. This problem is much less significant in fixed-grid WDM networks, where each connection is assigned a single wavelength. A number of defragmentation solutions have been proposed for flexible networks, broadly classified into proactive and reactive approaches. A proactive approach can be an RSA algorithm that avoids fragmentation by trying to leave spectrum voids that are more easily usable, while a reactive approach can be a special defragmentation algorithm that re-optimizes the network and frees spectrum for new connections.

## NETWORK PLANNING AND OPERATION TOOLS

Today a number of network planning and operation tools for fixed-grid WDM and IP networks are available. Most of these tools come from specific equipment vendors and are intended to

be used along with their products. Network operators and equipment vendors use these tools to meet service level agreements, achieve CAPEX savings, maximize network lifetime, and gain insight into the capabilities of their network.

Some of the most important functionalities these products provide are the following:

- Optimize routing and equipment placement to efficiently meet traffic demands
- Produce equipment configuration and equipment requirements
- Perform capacity planning, determining ways to expand the network (e.g., purchase links) to handle traffic growth
- Analyze the impact of failures and plan protection strategies to maximize resiliency
- Analyze and minimize equipment costs
- Evaluate various what-if networking scenarios, validating network changes or situations before deploying the production network
- Perform traffic analysis and engineering
- Visualize the above information

These commercial tools, some of which are presented in Fig. 3, offer functionality for the IP and optical WDM layers, most of the time separately. This is because changes in the two network layers occur at different timescales, as the WDM layer is very static and does not really need to be operated in real time. The introduction of flexible technologies will force tools to consider the new specifications, as the adaptability of the flexible paradigm brings the optical layer closer to the IP layer. For example, the BVT transponder and flex-grid switches would be reconfigured to accommodate that request in close to real time, so an operator no longer has

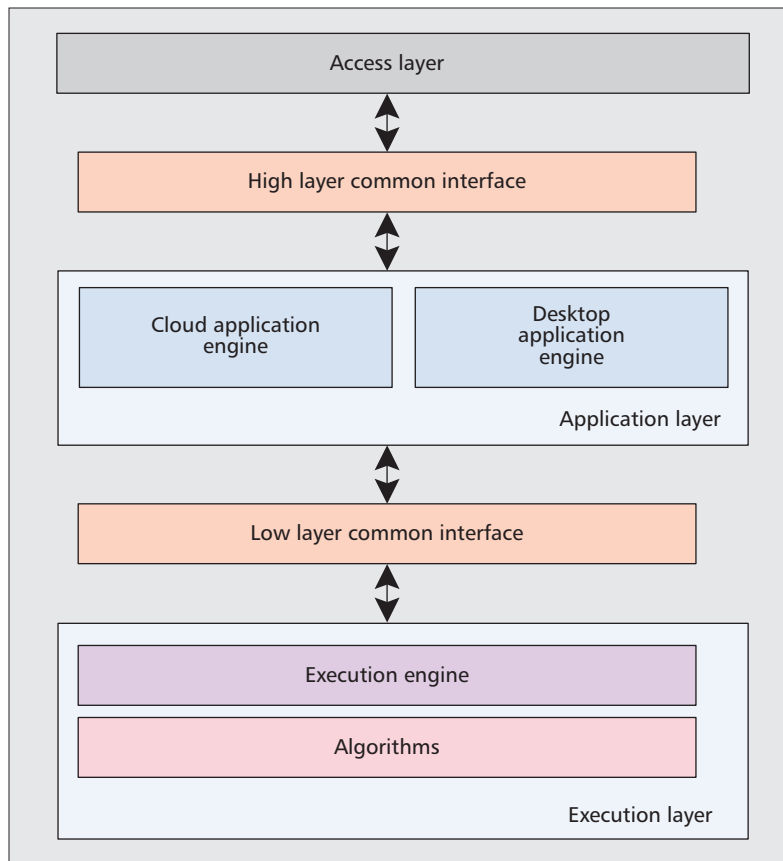


Figure 4. Mantis architecture main components.

to massively overprovision to accommodate possible fluctuations. It is thus evident that future tools will have to integrate these two layers more closely to achieve a more efficient resource utilization and save CAPEX/OPEX.

On a different note, the tools in Fig. 3 are standalone applications that might face scalability issues when considering huge what-if scenarios or controlling large multi-layer networks. Having applications operating in the cloud and harvesting its infinite computation power can be a solution to such issues.

#### MANTIS: A NETWORK PLANNING AND OPERATION TOOL FOR FLEXIBLE NETWORKS WITH CLOUD CAPABILITIES

Mantis is a network planning and operation tool for flexible and fixed-grid optical networks designed to include a variety of novel algorithms for planning and operation functions that have been developed over the past 5 years. Mantis architecture permits fast execution of the included algorithms, efficient usage of the computational resources, and deployment both as a desktop application and as a cloud service (SaaS).

In Mantis, users can define various parameters (e.g., network topology, traffic demands, equipment, device monetary and energy costs) and select among a large set of RWA and RSA algorithms, taking into account in both cases the PLIs. Algorithms evaluate future network plans and demands and report

detailed solutions including the required spectrum to serve the demands, the number, placement and configurations of transponders and regenerators, the total monetary and energy cost, etc.

#### MANTIS ARCHITECTURE

Mantis components are organized in three layers: the access, the application and the execution layer (Fig. 4). There are two common interfaces to provide loose coupling between the application and the other two layers, so that the same access and execution layers can be used whether Mantis is deployed as a desktop application or as a cloud service. The access layer handles the interaction with the users through a web-based interface and exposes a RESTful API. The execution layer consists of the execution engine and the library of available network planning and operation algorithms. The execution engine receives requests, for starting or terminating algorithms' executions, through the common interface from the application layer and is responsible for performing all the required actions, including the preparation of the execution environment, the monitoring of the execution progress and the handling of final results or possible failures. In the current version, the algorithms are written either in Cython or in C++ language and are accessed from the execution engine through a custom plug-in mechanism. This mechanism enables new algorithms to be easily added in the tool without any modification of the application layer and the execution engine.

The application layer orchestrates the execution of user requests. It is the only layer that differs between desktop application and cloud service deployment as there are different requirements and operations that should be performed in these different settings. In the former case, there is a single server that contains the desktop application engine and the execution layer. In the latter case the application layer implements the cloud application engine that handles the interaction with the cloud infrastructure. Mantis places one execution engine in every virtual computing node that runs a particular problem instance, and many virtual nodes can be placed in a physical machine in the cloud. The cloud application engine is responsible for the mapping of virtual nodes to physical machines and for adding and removing unused machines to the pool of Mantis resources, what we call the Mantis cloud scheduling policy. The cloud application engine is modular so as to support multiple cloud service providers with minimum effort and changes. In the current version, Mantis supports Amazon Web Services (AWS) and ~okeanos, GRNET's cloud service for the Greek Academic Community.

Figure 5 shows the total execution time of a variable number of standalone problem instances when Mantis operates in the desktop and in the cloud mode. We parallelize the execution of instances by allocating one instance per CPU core. The desktop application server run Ubuntu 12.04, and had 4 Gbytes RAM and 4 CPU cores, affording the ability to run 4 problem instances simultaneously. The cloud engine

scales dynamically, adding new machines to the pool when running out of cores. From Fig. 5 it is evident that the cloud deployment offers a high degree of parallelization, considerably decreasing the time required to run a set of experiments.

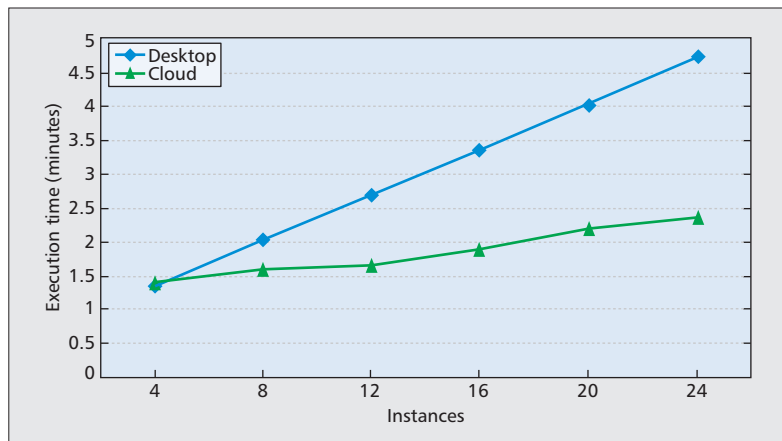
Mantis comes with a simple web-based user interface through which users access all its functionalities. Mantis user interface enables users to easily and graphically design network topologies, store, edit, and use them later. Users define their own traffic matrices, either graphically or by importing them from comma-separated values (CSV) files. Figures 6a and 6b show the available interface for network topologies and traffic demand creations, respectively. New configurations can be created for each algorithm by selecting a network topology, traffic demands and specifying all the other required parameters (Fig. 6c). Users can check the status of the running instances and have access to useful details (Fig. 6d). Furthermore, charts can be created by combining the results of completed instances. More details on Mantis can be found in [13].

### ALGORITHMS INCLUDED IN MANTIS

The current Mantis version includes network planning and operation algorithms for fixed and flexible optical networks that can be used for both transparent (without regenerators) and translucent (with regenerators) networks.

The IA-RSA (IA stands for Impairment Aware) algorithm [9] considers the planning problem of a flexible optical network under PLIs and addresses the problems identified earlier in a unified manner. It serves the demands for their requested rates by choosing the route, breaking the transmission into multiple connections, placing regenerators if needed, and allocating spectrum to them. Also implemented are two algorithms for planning mixed and single line rate WDM systems under physical layer impairments (IA-RWA-MLR and IA-RWA-SLR). The objective of all the aforementioned algorithms is to minimize a weighted combination of the cost of the transponders/regenerators and the spectrum that is used. Online (operation phase) versions of these algorithms are also available, where the output of the offline (planning phase) case is taken as input and a new demand is served in an incremental manner. Furthermore, the energy-aware (EA)-RWA algorithm of [12] is also included in Mantis, which aims at minimizing the energy consumed by the optical layer components when planning translucent WDM optical networks.

The algorithms included in Mantis are heuristic and utilize simulated annealing to obtain good solutions, while in the future we plan to connect Mantis with optimization solvers and incorporate algorithms implementing linear optimization formulations (i.e., LP and ILP). More algorithms that address dynamic network operation of flexible networks (discussed earlier), including dynamic spectrum expansion/contraction of existing connections, spectrum defragmentation, consideration of the IP layer, and others, are currently evaluated, and will be incorporated in future releases of Mantis.



**Figure 5.** Mantis execution time in desktop and cloud deployments for various number of instances.

### MANTIS PURPOSE

Our goal is to make Mantis a common benchmarking environment with social characteristics where researchers share topologies, traffic matrices, and CAPEX/OPEX parameters, and evaluate their algorithms under common conditions. In this way, Mantis could evolve into an online collaboration platform for optical network researchers, improving the comparability, quality, and reliability of the results presented in various research articles and projects. Toward a different direction, we are also working on connecting Mantis with the related management/control plane stacks of optical and IP equipment in order to make Mantis the logic of a complete network planning and operation tool that makes decisions and applies them to the underlying network (as shown in Fig. 1). A special case of that is to play the role of the path computation element (PCE) in a flexible optical network.

### CONCLUSIONS

Today a number of companies (equipment vendors, independent companies) offer planning tools for optical fixed-grid WDM and separate tools for IP networks. However, the introduction lately of flexible optical networking technologies raises the need for the development of new tools, such as:

- New algorithmic issues appear (routing and spectrum allocation, defragmentation).
- Existing ones change considerably (physical impairments, survivability).
- The optical network becomes more dynamic and comes closer to the IP layer, making cross-layer optimization important.

Mantis is one of the first tools implemented for planning flexible networks, and includes novel and efficient algorithms. Mantis modular architecture provides fast execution and efficient computation resource usage, and enables the deployment of the tool as both a desktop application and a cloud SaaS. Mantis' web-based user interface, along with its collaboration and social-related functionalities, could lead to an online collaboration platform for optical network research.

Our goal is to make Mantis a common benchmarking environment with social characteristics where researches share topologies, traffic matrices and CAPEX/OPEX parameters, and evaluate their algorithms under common conditions. In this way, Mantis could evolve to an online collaboration platform for optical network researches.

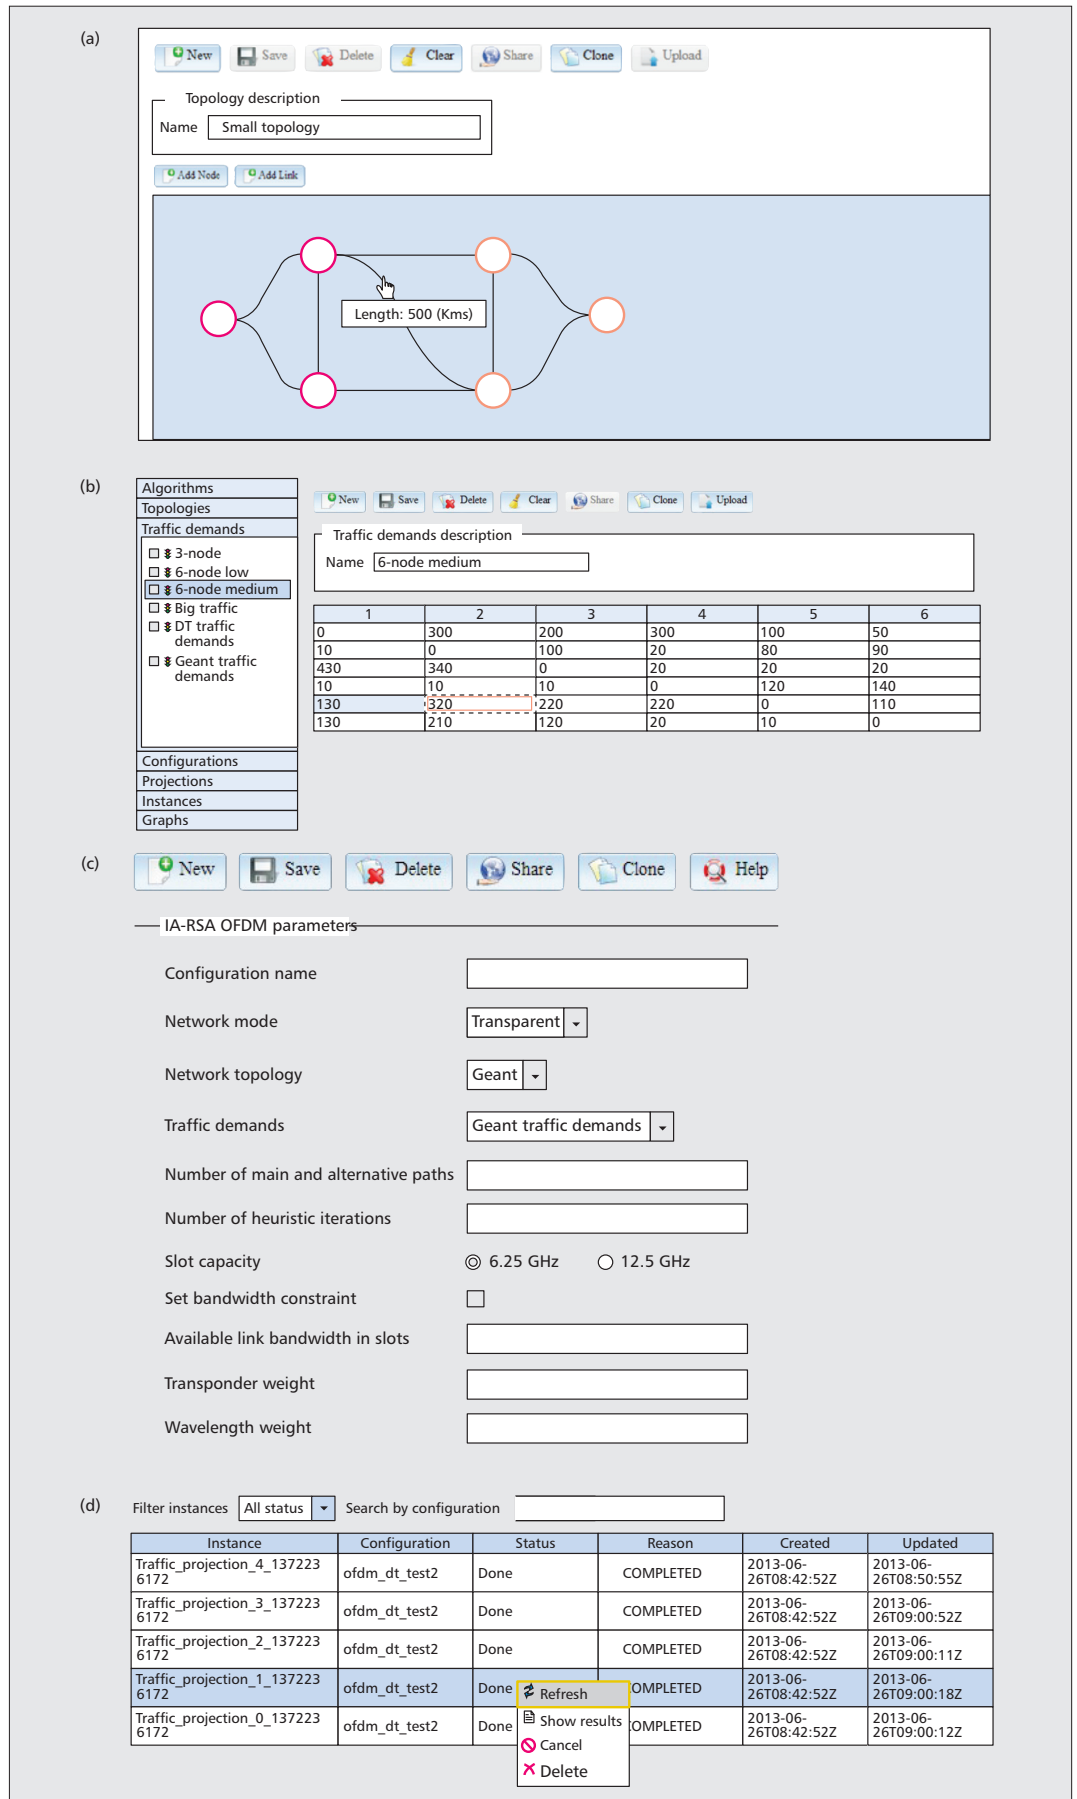


Figure 6. Creation of: a) network topology; b) traffic matrix in Gb/s; c) configuration; and d) display of information for different instances.



## ACKNOWLEDGMENTS

This work has been supported by the ICT IDEALIST project (grant agreement number 317999). P. Kokkinos is funded within the framework of the Action Supporting Postdoctoral Researchers of the Operational Program Education and Lifelong Learning (Action's Beneficiary: GSRT), co-financed by the European Social Fund (ESF) and the Greek State.

## REFERENCES

- [1] Cisco Visual Networking Index -Forecast and Methodology, 2012–2017.
- [2] T. Stern *et al.*, *Multiwavelength Optical Networks: Architectures, Design and Control*, Cambridge Univ. Press, 2nd ed., 2008.
- [3] O. Gerstel *et al.*, "Elastic Optical Networking: A New Dawn for the Optical Layer?," *IEEE Commun. Mag.*, 2012.
- [4] A. Autenrieth *et al.*, "Evaluation of Technology Options for Software-Defined Transceivers in Fixed WDM Grid versus Flexible WDM Grid Optical Transport Networks," *Photonics Network Symp.*, 2013.
- [5] FP7 IDEALIST project: <http://www.ict-idealists.eu/>
- [6] M. Jinno *et al.*, "Distance-Adaptive Spectrum Resource Allocation in Spectrum Sliced Elastic Optical Path Network," *IEEE Commun. Mag.*, vol. 48, no. 8, 2010.
- [7] A. Klekamp *et al.*, "Limits of Spectral Efficiency and Transmission Reach of Optical-OFDM Superchannels for Adaptive Networks," *IEEE Photonics Tech. Letters*, vol. 23, no. 20, 2011.
- [8] O. Rival, A. Morea, "Cost-Efficiency of Mixed 10-40-100 Gb/s Networks and Elastic Optical Networks," *OFC*, 2011.
- [9] K. Christodoulopoulos *et al.*, "Planning Flexible Network under Physical Layer Impairments," accepted, *IEEE/OSA J. Opt. Commun. and Net.*
- [10] K. Christodoulopoulos *et al.*, "Time-Varying Spectrum Allocation Policies in Flexible Optical Networks," *IEEE JSAC*, vol. 31, no. 1, 2013.
- [11] A. Patel *et al.*, "Defragmentation of Transparent Flexible Optical WDM (FWDM) Networks," *OFC*, 2011.

[12] K. Manousakis *et al.*, "Energy Efficient RWA Strategies for WDM Optical Networks," *IEEE/OSA J. Optical Commun. and Net.*, vol. 5, no. 4, 2013.

[13] Mantis: <http://www.mantis-tool.net/>

## BIOGRAPHIES

ARISTOTELIS KRETSIS graduated in 2006 from the University of Patras and was awarded his Diploma of Computer Engineer and Informatics. In 2009 he received his M.Sc. degree in computer science and engineering. He is currently pursuing a Ph.D. degree. His research interests are in the areas of grid computing, distributed computing, and software management tools.

KONSTANTINOS (KOSTAS) CHRISTODOULOPOULOS (kchristodou@gmail.com) is a research fellow at the Department of Computer Engineering and Informatics (CEID), University of Patras. He received his Diploma in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), his M.Sc. degree from Imperial College London, and his Ph.D. degree from CEID. He worked as a research fellow at CTVR, Trinity College Dublin, and as a contractor for IBM Research Ireland. His research interests are in the areas of algorithms and protocols for optical networks.

PANAGIOTIS KOKKINOS is a research fellow at NTUA. He received his Ph.D. in 2010 from the Computer Engineering and Informatics Department of University of Patras, Greece, in the field of optical grid networks. He also holds an M.Sc. degree (2006) and a Diploma (2003) from the same department. His current research activities are in the areas of distributed computing and networks.

EMMANOUEL (MANOS) VARVARIGOS received a Diploma in Electrical and Computer Engineering from NTUA in 1988, and M.S. and Ph.D. degrees in electrical engineering and computer science from MIT in 1990 and 1992. He is a professor in the Department of Computer Engineering and Informatics at the University of Patras, and the scientific director of the Network Technologies Sector of the Research Academic Computer Technology Institute.

*Today a number of companies (equipment vendors, independent companies) offer planning tools for optical fixed-grid WDM and separate tools for IP networks. However, the introduction lately of flexible optical networking technologies raises the need for the development of new tools.*