Net2Plan: An open-source network planning tool for bridging the gap between academia and industry

Pablo Pavon-Marino (*), Jose-Luis Izquierdo-Zaragoza

Universidad Politécnica de Cartagena (Spain)

{pablo.pavon, josel.izquierdo}@upct.es

(*) Corresponding author

Abstract

The plethora of network planning results published in top-ranked journals is a good sign of the success of the network planning research field. Unfortunately, it is often difficult for network carriers and ISPs to reproduce these investigations on their networks. This is partially because of the absence of a software planning tool meeting the requirements of industry and academia, which can make the adaptation and validation of planning algorithms less time consuming. We motivate how a paradigm shift to an open-source view of the network planning field emphasizes the power of distributed peer-review and transparency to create high-quality software at an accelerated pace and lower cost. Then, we present Net2Plan, an open-source Java-based software tool. Built on top of a technology-agnostic network representation, it automates the elaboration of performance evaluation tests for user-defined or built-in network design algorithms, network recovery schemes, connection-admission-control systems, or dynamic provisioning algorithms for time-varying traffic. The Net2Plan philosophy enforces code reutilization as an open repository of network planning resources. In this paper, a case study in a multilayer IP-over-WDM network is presented to illustrate the potential of Net2Plan. We cover standard CAPEX studies, and more advanced aspects such as a resilience analysis of the network under random independent failures and disaster scenarios and an energy efficiency assessment of "green" schemes that switch off parts of the network during low load periods. All the planning algorithms in this paper are publicly available on the Net2Plan website.

Keywords: Network planning tool, open-source software, technology-agnostic environment, IP networks, WDM networks

1. Introduction

Continuous advances in electronics and photonics produce a never-ending (re)evolution in the Telecom world, providing higher and higher capabilities to the network at a rapid pace. This challenges vendors and operators in their attempt to foresee which new technologies deserve resources and investments for deployment, which existing infrastructure to keep, and how to interoperate legacy and new elements during continuous migration processes. Backbone networks are a paramount example. Optical switching in optical wavelength division multiplexing (WDM) networks empowered by Optical Add/Drop Multiplexers (OADM) emerged in the 1990s as the enabling solution for addressing the so-called electronic bottleneck: the breach between the enormous transmission bandwidth of a fiber (tens of

Tbps) and the electronic packet processing limits of router line cards (tens/hundreds of Gbps). Backbone networks today are built according to the multilayer concept, where IP/MPLS traffic flows are routed over optically switched lightpaths of typical rates 10/40/100 Gbps, which are in its turn routed over the physical topology comprised of deployed fibers. Multilayer optical networks have stemmed the development of new cross-layer approaches to network planning, combining optimization of the optical and electronic layers. In parallel, there has been a growing interest in network features such as energy-efficiency, dynamic resource provisioning under varying traffic and network resilience under catastrophic failures. All these elements represent fundamental challenges for the network planning community.

Network planning tasks are assisted by software tools, so-called network planning tools. Different network planning tools can be found spanning a wide range of platforms, systems, languages, functionalities and applications. Some of them are oriented to the industry, whereas others are developed by academia for educational and research purposes. On the academia side, researchers investigating novel planning problems commonly need full control of the planning decisions, and develop their algorithms almost from scratch. In contrast, the market of commercial planning tools of IP/MPLS and optical networks is dominated by a set of third-party suites such as Riverbed OPNET NetOne, Cariden MATE Design (acquired by Cisco in 2012), RSoft MetroWAND or WANDL IP/MPLSView and NPAT (acquired by Juniper in 2014). All of them provide a complete set of features to design and analyze networks, without relying on a specific vendor. Simulation of several configuration scenarios, routing schemes, network recovery tests, or traffic load analysis, represent a small subset of those features. Some of them even provide capabilities to automate the configuration of network equipment of different vendors from the network plan.

In this context, the motivation of this work is our perception of a growing gap between the prolific research in network planning for multilayer networks, and the slow pace at which these network planning results are made available to the industry. In our opinion, this is partially caused by the absence of a planning tool that serves as a bridge between industry and academia, making the adaptation and validation of planning algorithms less resource consuming.

2. On the role of open-source network planning

As a research field, network planning is a recurrent hot topic. Year after year, R&D efforts are translated into a large number and variety of contributions. An illustrative example is the following: more than two-hundred works are indexed in IEEEXplore under terms “optical” and “network planning” in the period 2010-2012. In this context, research solutions are continuously undergoing major changes not only in terms of technologies considered, but also in methodologies. Nowadays, cost reduction, security, reliability and sustainability are some of the major drivers.

However, it is often difficult for network carriers and Internet service providers (ISPs) to apply in their networks the prospective studies and investigations found in the literature. On the one hand, network planning proposals from academia usually rely on own-developed software, and its source code is seldom provided (or even documented). Consequently, it is quite time-consuming for the industry to repeat, compare and reutilize these results. On the other hand, commercial tools have the natural incentive to provide support just for mature technologies and protocols for which there is a definite and large market, and have little or no flexibility to integrate new algorithms or test new technologies. Actually, they have evolved into “drag-and-drop” applications, with improved graphical interfaces to ease the production of “standard” planning studies in mature technologies with little effort, relying on non-disclosed heuristics to plan the network. As a side effect, despite of being a decisive element in network planning, the ability of these heuristics to find minimum cost/maximum performance solutions is assumed as a leap of faith. Actually, we believe that commercial tools do not compete in practice in this aspect. As an example, institutions holding academic licenses are sometimes forbidden to publish any study comparing the quality of the results of the (not disclosed) algorithms [1].

We believe that open-source software can have an important role in this context to meet the requirements of both academia and industry. In the real world, no business is static and software changes to suit users’ needs are required, and arguably, network planning is not an exception. In general, open-source software is closer to what users actually demand, since those users are able to tweak it to suit their needs. In addition, open-source could be deemed as peer-reviewed, which leads to more reliable and robust software. For example, Internet infrastructure is largely composed of both open-source programs (i.e. DNS) and languages (i.e. HTML).

From our point of view, an open-source view of the network planning world can bring advantages to industry and academia. First, it encourages a large market of early adopters (compared to that for commercial products) who actively help to debug the software. Consequently, open-source software becomes highly robust at a surprisingly early stage of its development. Second, it stimulates the building of code repositories, so that users are able to share their knowledge (i.e. code contributions), and accelerate collaboration cycles (i.e. technology transfer). Third, users are not limited of the vendor's vision, priorities and timetable. Integration of new functionalities it is simply a matter of deploying them. Finally, open-source projects translate into substantial cost savings, since typically do not have a per-seat license model.

Following this paradigm, we developed Net2Plan [2]. Net2Plan is an open-source (multilayer) network planning tool, and a repository of planning resources. Net2Plan is not constrained to any specific network technology, and is adaptable to any of them. It allows users to rapid-prototype their own algorithms, or use the provided built-in ones. Users can evaluate their designs using either automatic report generation or post-analysis tools for network resilience, connection-admission-control and time-varying traffic resource allocation. Net2Plan enforces code reusability in two ways. The open-source nature enables code reutilization and public validation. The technology-agnostic nature permits reusing algorithms and studies for similar problems appearing in different network technologies. We believe that identifying these similarities permits incorporating well-known results in “old” technologies, to improve and speed-up the development of algorithms for “new” technologies.

3. Net2Plan

Net2Plan is an open-source Java-based software, publicly and freely available to download from its website [2]. It is licensed under the GNU Lesser General Public License (LGPL).

Net2Plan has its origins in September 2011, as a resource for network planning courses at Technical University of Cartagena. As a sign of its stability, in 2013 it was used during more than 50 hours of lab work in two undergraduate courses, summing more than 150 students.

Net2Plan is designed with the aim to overcome the barriers imposed by existing network planning tools for two main reasons: (i) users are not limited to execute built-in algorithms, but also can integrate their own algorithms, applicable to any network instance, as Java classes implementing particular interfaces (see Fig. 1), and (ii) Net2Plan defines a network representation, so-called network plan, based on abstract concepts such as nodes, links, traffic demands, routes, protection segments, shared-risk groups and network layers. In addition, technology-specific information can be introduced via user-defined attributes attached to nodes, links, etc. in the network plan. The combination of a technology-agnostic substrate and technology-related attributes provides the required flexibility to model any network technology within Net2Plan.

Net2Plan provides both a graphical user interface (GUI) and a command-line interface (CLI). In either mode, the current version of Net2Plan (0.2.3, March 2014) includes six tools:

- Offline network design: Targeted to execute offline (multilayer) planning algorithms (see Fig. 2). Algorithms based on constrained optimization formulations (e.g. integer linear programs, or ILPs) can be fast-prototyped using the open-source Java Optimization Modeler library (JOM [3]), to interface to a number of external solvers such as GPLK, CPLEX or IPOPT. Algorithms for offline network design should implement *IAlgorithm* interface (see Fig. 1).
- Traffic matrix generation: Assists users in the process of generating and normalizing traffic matrices.
- Resilience simulation: Permits evaluating the availability performance of online protection and restoration algorithms in the network. Provisioning algorithms that react to network failures and reparations, should implement the *IProvisioningAlgorithm* interface (see Fig. 1). Modules that generate failure and reparation events to be consumed by provisioning algorithms should implement the *IResilienceEventGenerator* interface (see Fig. 1).
- Connection-admission-control (CAC) simulation: Targeted to analyze the blocking performance of online provisioning schemes that allocate resources to incoming connections (i.e. on-demand virtual circuit requests). CAC algorithms that react to connection requests, should implement the *ICACAlgorithm* interface (see Fig. 1). Modules that generate connection requests to be consumed by CAC algorithms should implement the *IConnectionEventGenerator* interface (see Fig. 1).
- Time-varying traffic simulation: Permits evaluating the performances of online algorithms that react to traffic variations (e.g. traffic rerouting schemes). Allocation algorithms that react to variations in the traffic demands (e.g. used in the time-varying traffic simulation tool), should implement the *ITrafficAllocationAlgorithm* interface (see Fig. 1). Modules that generate time-varying traffic events to be consumed by provisioning algorithms should implement the *ITrafficGenerator* interface (see Fig. 1).
- Reporting: Net2Plan permits the generation of built-in or user-defined reports, from any network design. The report generation tool is integrated within all the previous

functionalities, so that it is possible to create reports collecting performance measures in any of these aspects. Reports should implement the *IReport* interface (see Fig. 1).

We recall that every algorithm, report, event generator, and provisioning/allocation algorithm in Net2Plan can be either built-in (from the repository [2]) or user-made. We refer the reader to [2] for fully-detailed information regarding Net2Plan functionalities.

3.1. Benefits, shortcomings and roadmap

To the best of the authors' knowledge, Net2Plan is the only open-source network planning tool with the complete set of features described, built up over a technology-agnostic philosophy, and open interfaces to develop technology-specific extensions. Net2Plan enforces a collaborative framework to share research works and speed-up their application to operational networks. Nonetheless, industry might be reluctant to use Net2Plan, or any other open-source tool, for different reasons. Two main concerns are identified: confidentiality and technical support. Our licensing model (LGPL) addresses the first issue, stimulating technology-transfer. In LGPL software, commercial modifications and algorithms in Net2Plan are possible, and not forced to be open-source. Regarding to technical support, many open-source companies follow a dual-licensing model, where commercial support contracts are offered to customers. Instead, we believe that community participation is a critical success factor for any open-source project, and Net2Plan may evolve into a collaboration platform, connecting people and information within an online ecosystem, for example, including a forum and a bug tracker.

On the teaching and training side, Net2Plan promotes self-learning, and provides an ever-growing documentation, including e.g. video tutorials or teaching materials [2]. As a teaching resource, students and practitioners can use Net2Plan to examine the planning of particular technologies. Net2Plan is also a valuable resource for courses concentrating on network optimization concepts and transversal network planning skills, avoiding an in-depth study of technologies apparently different. In fact, this training strategy was already proposed by Doshi *et al.* [4] more than a decade ago.

However, compared to commercial tools, Net2Plan still has a long road ahead and some lacks should be addressed in the future. On the one hand, all commercial tools provide a powerful GUI for users' interaction. Features such as multilayer visualization or chart generation, combining data from multiple sources, are to be implemented. On the other hand, emergence of software defined networking (SDN) paradigm and initiatives like OpenDaylight [5] may represent a clear opportunity to develop extensions to evolve into an in-operation network planning tool, providing interaction with equipment and network management systems (NMS).

4. Case study: Designing an IP-over-WDM network

In this section, we present a complete real-world case study based on the planning and evaluation of a multilayer IP-over-WDM network. The planning results provided cover both classical problems for which state-of-the-art algorithms have been used (or reutilized with some variations), together with quite novel algorithms and recent studies. All the algorithms are publicly available in [2]. In our opinion, the main contribution is not the scientific value of algorithms and results, but to illustrate the advantages that open-source and technology-agnostic tools like Net2Plan could provide to the network planning field.

In our case study, we focus on a vertically integrated operator owning an IP-over-WDM multilayer network, using as a model the well-known 14-node 42-link NSFNET network and a reference IP end-to-end traffic matrix [6]. In this context, IP routers with grooming capabilities are connected via a virtual topology of 40 Gbps lightpaths, optically switched by OADM equipment. To guarantee network survivability under single fiber link failures, each lightpath is realized as two link-disjoint routes (primary and backup), according to a 1+1 protection scheme, transparently to the IP layer. The maximum reach of a lightpath (primary or backup) is limited to 2800 km, which is consistent with the transparent reach of modern coherent optical modulations [7]. Lightpaths exceeding this distance need to go through regenerators located at nodes, which also permit wavelength conversion. Finally, the IP traffic engineering is governed by the OSPF protocol with statically fixed weights for each link (1+1 lightpath pair).

4.1 Modeling an IP-over-WDM network within Net2Plan

IP-over-WDM networks are modeled using a two-layer network representation. The (upper) IP layer has a set of electronic traffic demands defined by the IP end-to-end traffic matrix. This traffic is routed over a set of links, one per lightpath. The multilayer concept makes each link in the upper layer (lightpath) become a demand at the lower (optical) layer. Then, lightpaths are routed at the lower layer over the given set of fibers (links), being their route the routing and wavelength assignment (RWA). The IP-over-WDM technology-specific attributes in this network representation are as follows. For fiber links (links at the lower layer), we introduce the attribute “numWavelengths” to represent the number of WDM channels available (in this work, we assume 80 channels per fiber); for working and protection lightpaths (represented by routes and protection segments, respectively), we introduce the attributes “seqWavelenghts”, a vector with the wavelengths used in each traversed fiber, and “seqRegenerators” a vector indicating with 1/0 values whether or not a regenerator (or wavelength converter) is installed in each traversed node. Lightpaths have two attributes: (i) “linkWeight” as the OSPF weight at the IP layer, (ii) “capacity” given by its nominal rate (40 Gbps). We use the Net2Plan libraries *IPUtils* and *WDMUtils* to check the validity of the attributes (i.e. two lightpaths cannot use the same wavelength in the same fiber). These libraries have been developed to ease the fast-prototyping of algorithms in IP-over-WDM networks.

4.2 Offline network planning: CAPEX estimations

In our case study, the network operator has a traffic demand of a volume proportional to the reference traffic matrix [6]. In year 2014, the total traffic volume will be 1 Tbps, and forecasts estimate a compound annual growth rate (CAGR) of 20%. In this section, we describe a multilayer network design algorithm (algorithm NDA-1 in Table 1) devised to plan and allocate the resources at the IP and optical layers for a given year, minimizing the capital expenditures (CAPEX) cost of the network. The CAPEX cost considered is composed of a fixed cost per lightpath (given by the cost of the transponder, the slot card and short-reach interface in the routers), and a cost per optical regenerator equipment, being 60.68 and 5.17 monetary units, respectively (data obtained from [8]). The design is considered valid if two extra constraints are met: lightpath utilization must be below 50%, and IP traffic cannot suffer an end-to-end propagation delay higher than 50 ms (a typical value for maximum inter-PoP delay in the continental USA [9]). In NSFNET, the distance between the two farthest nodes is equal to 4500 km, which is roughly equivalent to 22.5 ms.

To the best of the authors' knowledge, the multilayer problem variant described has not been addressed yet in the literature. Still, it has been selected on purpose as a combination and variation of several previously addressed complex planning problems. In this respect, what follows is a high-level description of the planning algorithms devised, emphasizing how existing algorithms have been reused, modified and chained in Net2Plan. Full details together with the source code can be found in the examples repository in [2].

Algorithm NDA-1 decomposes the problem into two sub-problems addressed by two algorithms NDA-1.1 and NDA-1.2. First, NDA-1.1 is an RWA algorithm to find a 1+1 protected full-mesh virtual topology fulfilling the aforementioned design constraints. The algorithm is based on an ILP formulation, modeled with the JOM library [3], which optimally minimizes CAPEX. Then, algorithm NDA-1.2 is used to modify the previous design, adding traffic grooming to the picture by varying the OSPF weights, and potentially eliminating some lightpaths. The basic idea is to find a valid OSPF routing, which has feasible link utilizations and end-to-end delays, while iteratively lightpaths are removed, in descending order of cost, from the solution coming from NDA-1.1.

Determination of the optimal setting of OSPF weights is a classical traffic engineering problem, and several algorithms have been proposed to address it. Algorithm NDA-1.2 reutilizes a Java implementation of the state-of-the-art IGP-WO algorithm [10], based on the tabu-search meta-heuristic. The worst lightpath utilization, plus a penalization when end-to-end delay constraints are violated, is used as the new fitness function.

The multilayer algorithm developed allows creating CAPEX forecasts for the coming years, as shown in Fig. 3. Our algorithm could not find solutions carrying 100% of the traffic after year 2025, determining that a major network upgrade (i.e. light up more fibers) would be required at that time to satisfy the expected traffic demand.

Table 1. Java classes used in the case study. All the code is included in the Net2Plan repository and/or shipped with Net2Plan release.

Stage	Code	Description	Java class
Network design	NDA-1	IP over WDM multilayer algorithm	TCFA_IPoverWDM_OSPFAndVTDE2eDelayLimit
	NDA-1.1	RWA algorithm with regenerator placement	CFA_WDM_RWARPP
	NDA-1.2	IGP-WO algorithm for setting the OSPF weights enforcing traffic grooming	FA_OSPF_IGPWOVariation
Network resilience analysis	R-1	Report Availability	Report_availability
	R-2	Report Disaster Vulnerability	Report_disasterVulnerability
	FRGA-1	Failure/Reparation event generator	NRSim_EG_exponentialSRGFailureGenerator
	NRA-1	Provisioning algorithm (Protection)	NRSim_AA_genericProtectionSegmentAlgorithm
	NRA-2, NRA-3	Provisioning algorithm (Restoration)	NRSim_AA_WDM_pathRestoration
Energy-efficiency analysis	TVGA-1	Time-varying traffic generation algorithm	TVSim_EG_activityFunction
	TVAA-1	Time-varying traffic allocation algorithm	TVSim_AA_OSPF_switchOffLinkFixedWeight

4.3 A study of the network availability versus cost for different recovery systems

As a starting point, the operator considered the 1+1 scheme as a simpler and faster recovery system. However, he is willing to quantify if more flexible restoration schemes can improve cost and availability performances. Now, the operator is interested in evaluating three options for recovering lightpaths from failures: the 1+1 lightpath protection scheme already described, and lightpath restoration in two modes: path mode and sub-path mode [11]. Under restoration mode, each lightpath only reserves a primary route during set-up, ignoring the original backup 1+1 information. Upon failures, backup routes will be searched for by an i.e. GMPLS control plane. In the path mode, the restoration route is the shortest one (in km) with at least one idle channel in each traversed fiber. Upon route selection, wavelengths and regenerators are jointly assigned. In the sub-path mode, restoration is initiated by the node right before the failing link, following a fast-reroute scheme similar to the one employed in MPLS networks. In any restoration mode, when the primary route becomes active again, the lightpath route is reverted.

The operator is interested in assessing the three alternatives in two failure scenarios: (i) a standard scenario where failures can occur randomly over bidirectional fiber ducts, and (ii) a catastrophic multiple-failure scenario, where geographically close links and nodes suffer a simultaneous breakdown, i.e. because of natural disasters.

Independent failure scenario

Net2Plan resilience tools and reports are based on the *shared risk group* (SRG) concept. A SRG is a set of links and/or nodes assumed to share a potential cause of malfunction. Users can define the SRGs in the network, according to their identification of the resources that are subject to fail. For each SRG, the mean time to fail (MTTF) and mean time to repair (MTTR) information is provided. In our case study, one SRG is defined for each bidirectional fiber link. We use conservative standard MTTF and MTTR figures (8748 and 12 hours, respectively) which correspond to an average of one failure per year, with 12 hours of reparation time.

The operator is interested in evaluating the trade-off between performance and cost for the three recovery systems. On the performance side, we use as a figure the worst lightpath availability, where availability is defined as the fraction of time in which a resource is operative. On the cost side, we assume the 1+1 protection scheme requires twice the lightpath equipment, without considering regenerators, since backup routes are known and resources are allocated in advance, allowing negligible switching times. In contrast, restoration schemes compute backup routes according to the network state, and resources can be shared among lightpaths. Here, the operator is forced to place a sufficient number of regenerators at the nodes so that restoration routes can, with a sufficiently large probability, find available regenerators if needed.

We use two different Net2Plan functionalities to explore the performance versus cost trade-off. First, the *resilience simulation tool* is used to run an event-driven simulation of the network operation, where failure and reparation events randomly occur in the SRGs according to the MTTF/MTTR figures defined (failure/repair generator algorithm FRGA-1 in Table 1). The tool automates this type of test, for any recovery scheme. The recovery schemes have been implemented in the algorithms indexed as NRA-1 (1+1 protection), NRA-2 (path restoration) and NRA-3 (sub-path restoration) in Table 1. In our case, we simulated 1 million failure/reparation events, with a transitory period of 100000 events. The tool collects several metrics, including per-demand (i.e. lightpath in our case) availability. On the other hand, technology-dependent metrics are collected within NRA-2 and NRA-3. In particular, we are interested in the maximum number of regenerators allocated into each node throughout the simulation. We assume that this is the number of regenerators required to have a negligible probability of regenerator exhaustion.

Besides the resilience simulation tool, Net2Plan offers an *availability report* (R-1 in Table 1), as a second functionality suitable for our case. Given a network design (i.e. optical layer) and a recovery scheme (i.e. NRA-1, NRA-2 and NRA-3), several availability metrics are estimated. The report enumerates all possible single and double SRG failures (all other failure states, i.e. triple or quadruple, are assumed not to occur), computes their probability of occurrence, and obtains the resulting network state applying the given resilience algorithm.

Results are shown in Fig. 3 and Table 2. In Fig. 3, we observe that the cost of the two lightpath restoration modes is quite similar, being the sub-path scheme slightly more expensive. On the contrary, the difference with the 1+1 protection scheme is significant. Since it requires twice the lightpath equipment, which is the main contribution to the total cost, this alternative is the costly option, neglecting the discrepancy in number of regenerators.

Table 2 illustrates the worst-case lightpath availability performance for three different years for all three recovery schemes. As can be seen, five-nine availability (~5 minutes downtime per year) is provided by the two restoration schemes. Conversely, the 1+1 protection scheme has significantly worse availability, since it cannot recover lightpaths suffering from failures affecting both their primary and backup routes. As a conclusion, while the 1+1 protection scheme is a simpler and faster recovery system, lightpath restoration schemes are desirable since they are able to improve both cost and availability performances.

Table 2. Resilience metrics for different failure scenarios

Metric	Recovery scheme		Year		
			2014	2020	2025
Worst-case lightpath availability under the bidirectional fiber duct model	1+1	Report	99.724%	99.721%	99.724%
		Simulation	99.721%	99.714%	99.716%
	Path restoration	Report	99.999%	99.999%	99.999%
		Simulation	99.999%	99.999%	99.999%
	Sub-path restoration	Report	99.999%	99.999%	99.999%

		Simulation	99.999%	99.999%	99.999%
Average surviving traffic under the disaster failure model	1+1		80.60%	81.12%	82.09%
	Path restoration		85.71%	85.71%	85.71%
	Sub-path restoration		85.71%	85.71%	85.71%

Tests for multiple (disaster) failures

A growing interest exists in the networking community to assess the network vulnerability under unfortunate disasters provoking multiple failures. For these cases, complete network recovery is considered an unreachable target, and the interest shifts to evaluate which recovery schemes allows a larger portion of the traffic to survive. Net2Plan offers the disaster network vulnerability built-in report (R-2 in Table 1), suitable for this purpose. The set of multiple failure scenarios is defined via SRGs. Then, the report estimates, for a given recovery scheme, the fraction of network traffic that survives in each case. In our case study, the operator uses the information in [12] to identify a set of 14 potential node-centric multiple-failures (SRGs) that may occur, which also may include some geographically close links. Table 2 shows, for each recovery scheme, the fraction of surviving traffic averaged among the possible disaster-failures considered. As can be seen, restoration schemes show their clear strength in this aspect, thanks to their flexibility to decide the backup routes upon failure detection.

4.4. Energy-efficiency analysis

It is well-known that traffic in backbone networks fluctuates following a daily pattern [13]. Peak traffic is observed during working hours, traffic volumes drop at early-morning, evening and weekends. As the final stage of the study, the operator is interested in a prospective evaluation of the potential power-consumption savings obtained if these traffic patterns were exploited, by selectively switching to sleeping state low-loaded lightpaths, letting OSPF automatically reroute the traffic.

Several techniques have been proposed in recent years under the umbrella of the so-called *green networking*. Here, the operator considers the proposal in [13], based on a network controller unit (NCU) that uses traffic information to decide every 5 minutes which lightpaths to put into sleep mode. OSPF weights are never modified in order to maximize routing stability, and the maximum lightpath utilization is limited to 80%. Note that this threshold is less conservative than the 50% utilization figure used for static planning, but it may be acceptable for dynamic scenarios.

The *Time-varying traffic simulation* tool is used in this section to perform the energy-efficiency study. This tool runs a simulation where traffic varies according to a traffic generation algorithm, and network reactions are planned by an allocation algorithm. In our case, algorithm TVGA-1 governs the traffic generation, and TVAA-1 the allocation (see Table 1). In TVGA-1, end-to-end traffic depends on the *activity factor* of each node (considering the node time zone) like in [14]. Early-morning and evening traffic is 30% of the peak values, and on weekends the peak value is 50% of that on working days. A Gaussian randomness is added to each demand with a typical deviation of 5% of its traffic. The TVAA-1 algorithm allocates resources according to the proposal in [13] described above.

Fig. 4 plots the comparative power consumption figures obtained by applying or not the "green mode" in the network with 1+1 protection scheme, simulated for one month of real time. Power consumption is assumed to be independent of the traffic load, but negligible for devices in "green mode", and is computed considering only consumption of lightpath equipment (slot cards, short-reach interfaces, transponders) and regenerators, according to the model in [8] (948 and 100 W, respectively). Results show that roughly 50% power savings are achieved at every load (e.g. a reduction of 485 kW would be expected in year 2025).

5. Conclusions

In this paper, we motivate how open-source network planning tools would help to translate, at an accelerated pace and lower cost, research efforts from academia into real-world implementations by the industry. Then, we present Net2Plan, an open-source network planning tool. For researchers, we see Net2Plan as a valuable resource for easing the development and evaluation of novel schemes for offline network design, network recovery, connection-admission-control and dynamic network provisioning. Industry and academia can also benefit from the rapid-prototyping of planning algorithms that Net2Plan enforces. In addition, the Net2Plan philosophy promotes the creation of open repositories containing planning results that can strengthen the link between academia and industry. Our experience using Net2Plan as a teaching and training resource is highly positive. Teaching materials specifically devoted for network planning courses can be found in [2], together with a repository of algorithms and network planning resources. The paper includes a complete case study for a multilayer IP over WDM network, to show the functionalities of Net2Plan and its flexibility to model complex network technologies. Net2Plan has been successfully used for researching in other technologies like online provisioning algorithms for flexi-grid optical networks (see [15] and Net2Plan repository [2]). The application of Net2Plan to highly-dynamic scenarios like wireless sensor networks is considered a promising line of work.

Acknowledgements

This research was partially supported by the Spanish project grant TEC2010-21405-C02-02/TCM (CALM) and the FPU predoctoral fellowship program of the Spanish Ministry of Education, Culture and Sport (reference no. FPU12/04571). It was also developed in the framework of project "Programa de Ayudas a Grupos de Excelencia de la Región de Murcia" funded by F. Séneca (Plan Regional de Ciencia y Tecnología 2007/2010).

References

[1] Riverbed Technology, "University Research Program Requirements" [Online]. Available: http://media-cms.riverbed.com/documents/University+Program+for+Research+Requirements.pdf [Last Accessed: April 2014]

[2] Net2Plan – The open-source network planner [Online]. Available: http://www.net2plan.com/ [Last accessed: April 2014]

[3] JOM – Java Optimization Modeler [Online]. Available: http://ait.upct.es/~ppavon/jom/ [Last accessed: April 2014]

[4] B. Doshi, P. Harshavardhana, "Broadband network infrastructure of the future: roles of network design tools in technology deployment strategies," *IEEE Communications Magazine*, vol. 36, no. 5, pp. 60-71, May 1998.

[5] OpenDaylight Project [Online]. Available: http://www.opendaylight.org/ [Last accessed: April 2014]

[6] R. Ramaswami, K.N. Sivarajan, "Design of Logical Topologies for Wavelength-Routed Optical Networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 840-851, June 1996.

[7] A. Nag, M. Tornatore, B. Mukherjee, "On the effect of channel spacing, launch power, and regenerator placement on the design of mixed-line-rate optical networks", *Optical Switching and Networking*, vol. 10, no. 4, pp. 301–311, November 2013.

[8] A. Bianco, E. Bonetto, F. Musumeci, A. Pattavina, M. Tornatore, "CapEx/OpEx Evaluation of Circuit vs Packet Switched Optical Networks," in *Proceedings of the 17th International Conference on Optical Networking Design and Modeling* (ONDM 2013), April 2013.

[9] A. Nucci, K. Papagiannaki, *Design, measurement and management of large-scale IP networks*, Cambridge University Press, 2009.

[10] B. Fortz, M. Thorup, "Internet traffic engineering by optimizing OSPF weights", in *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies* (INFOCOM 2000), March 2000.

[11] J. Wang, L. Sahasrabuddhe, B. Mukherjee, "Path vs. Subpath vs. Link Restoration for Fault Management in IP-over-WDM Networks: Performance Comparisons Using GMPLS Control Signaling," *IEEE Communications Magazine*, vol. 40, no.11, pp. 80-87, November 2002.

[12] M. Farhan-Habib, M. Tornatore, M. De-Leenheer, F. Dikbiyik, B. Mukherjee, "Design of Disaster-Resilient Optical Datacenter Networks," *Journal of Lightwave Technology*, vol. 30, no. 16, August 2012.

[13] R. Bolla, R. Bruschi, A. Cianfrani, M. Listanti, "Enabling Backbone Networks to Sleep," *IEEE Network*, vol. 25, no. 2, pp. 26-31, March-April 2011.

[14] J. Milbrandt, M. Menth, S. Kopf, "Adaptive Bandwidth Allocation: Impact of Traffic Demand Models for Wide Area Networks," University of Würzburg, Germany, Report No. 363, June 2005.

[15] J.L. Izquierdo-Zaragoza, P. Pavon-Marino, M.V. Bueno-Delgado, "Distance-adaptive online RSA algorithms for heterogeneous flex-grid networks," in *Proceedings of the 18th International Conference on Optical Network Design and Modeling* (ONDM 2014), May 2014.

Biographies

Pablo Pavon-Marino (pablo.pavon@upct.es) received the M.Sc. degree in Telecommunication Engineering in 1999 from the University of Vigo, Spain. In 2000, he joined the Universidad Politécnica de Cartagena, Spain, where he is Associate Professor in the Department of Information Technologies and Communications. He received the Ph.D. degree from this University in 2004, and a M.Sc. degree in Mathematics in 2010. His research interests include the planning and optimization of communication networks.

Jose-Luis Izquierdo-Zaragoza (josel.izquierdo@upct.es) received his B.Sc. in Telecommunications and M.Sc. in Information Technologies and Communications in 2010 and 2011, respectively, from the Universidad Politécnica de Cartagena, Spain, where he is currently working toward the Ph.D. degree in the Department of Information Technologies and Communications. His research interests include planning and performance evaluation of optical networks.

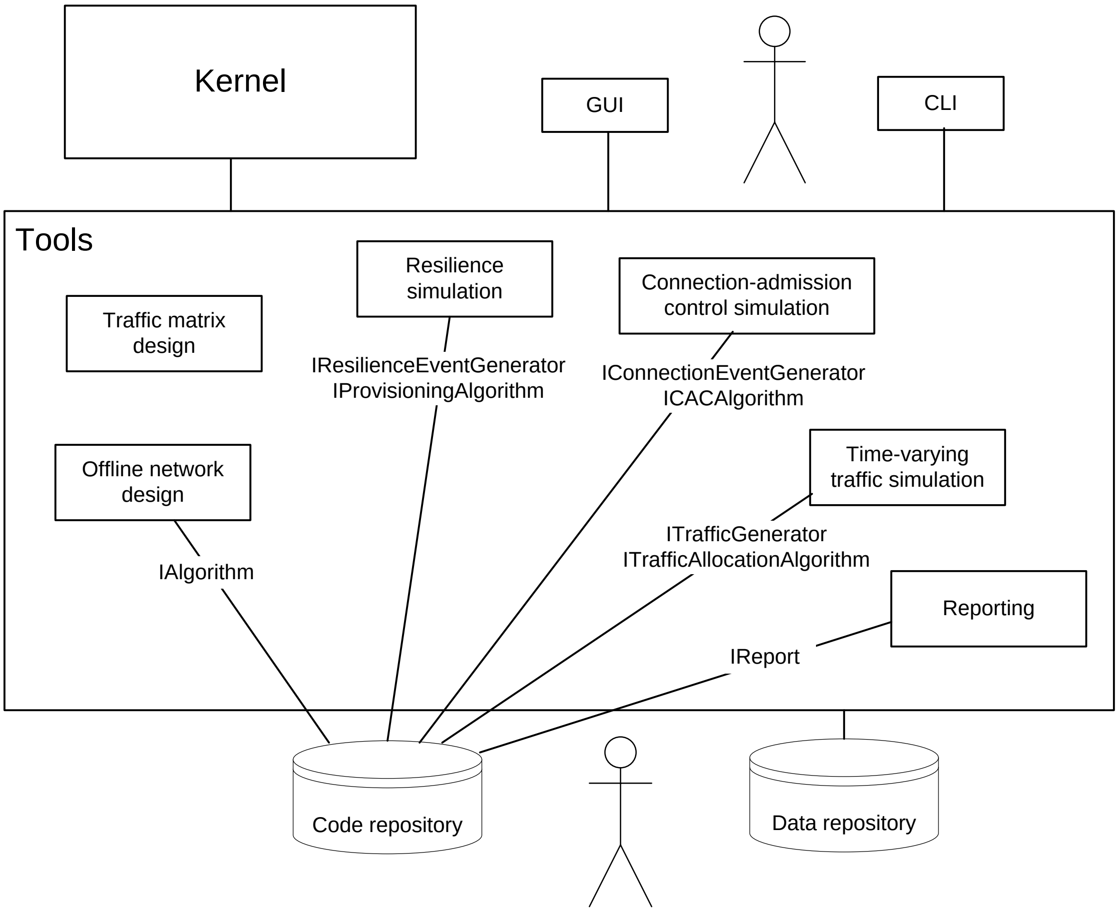

Figure 1. Architecture of Net2Plan

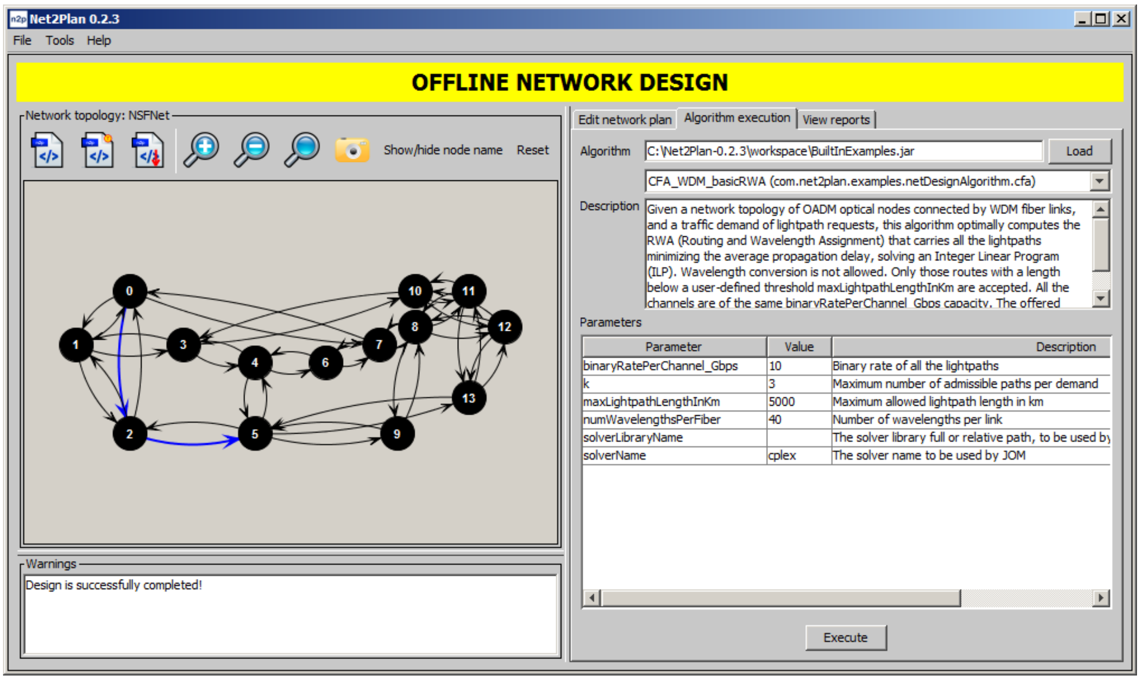

Figure 2. Offline network design tool

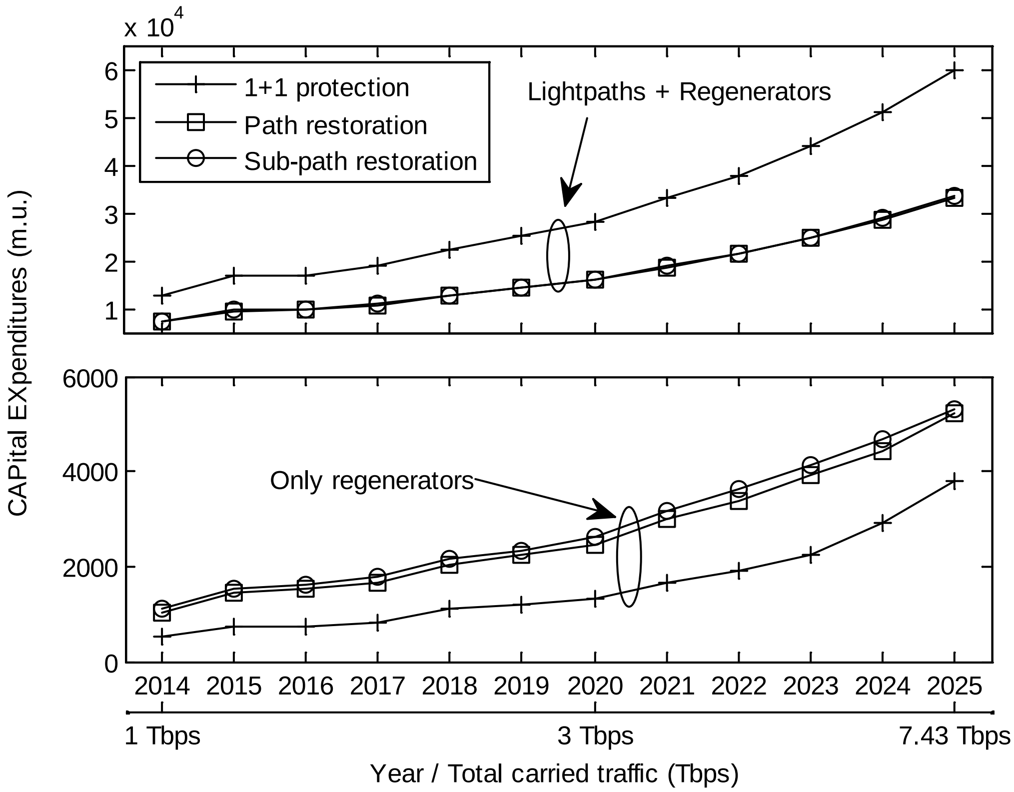

Figure 3. CAPital EXpenditures (CAPEX) for different resilience approaches

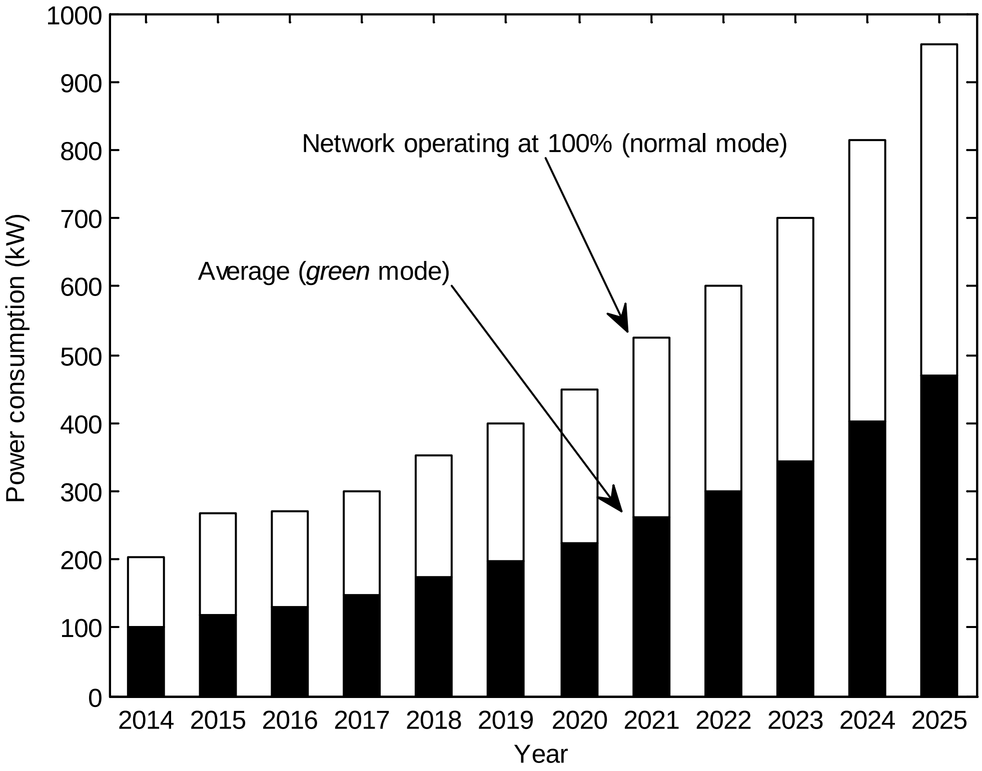

Figure 4. Power-consumption comparison achieved with and without using an energy-efficient traffic engineering algorithm