Performance evaluation of a Bayesian decisor in a multi-hop IP over WDM network scenario

Víctor López, José Alberto Hernández, Javier Aracil
Universidad Autónoma de Madrid, Spain
C/ Francisco Tomás y Valiente 11, E-28049 Madrid, Spain

Óscar González de Dios, Juan P. Fernández Palacios
Telefónica I+D, Spain
C/ Emilio Vargas, 6, E-28043 Madrid, Spain

Abstract—**Network operators have understood the importance of migrating their backbone networks to IP over WDM architectures, whereby an underlying optical infrastructure can automatically set up and tear down direct optical connections (lightpaths), yet keeping electronic IP routers on top of it. In such multi-layer networks, it is necessary to efficiently combine the resources available in both electronic and optical layers, providing the necessary Quality of Service (QoS) to end-users at the minimum possible cost.**

This work defines a multi-layer Bayesian decisor in a multi-hop scenario which finds a compromise between the utility perceived by the users in terms of delay and the utilisation costs of the optical and electronic resources. The mathematical formulation of such a Bayesian decisor is formulated and its behaviour is further analysed for different configurations in realistic scenarios. Its behaviour shows that the algorithm is capable of adapting its decision according to traffic characteristics, while utilising only the necessary optical and electronic resources.

I. INTRODUCTION

The IP over WDM paradigm appears as the most promising technology for future backbone architectures. Essentially, these architectures are based on high-capacity IP routers interconnected by high-speed point-to-point links of the underlying WDM network. Remark that such an underlying optical infrastructure provides potential bandwidths of Terabits per second, since each optical fibre link offers tens or hundreds of wavelengths at several Gigabits per second of speed, thanks to WDM. Such a technology is expected to meet the ever-increasing user demands for bandwidth capacity, at least in the near future.

With the advent of specialised optical hardware, such as Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Optical Cross Connects (OXC), backbone nodes have become “multi-layer” capable. In other words, such multilayer nodes are able to either transmit their traffic demands using the IP layer of intermediate nodes (OEO switching) or directly through an all-optical end-to-end lightpath (OOO switching), thanks to the unified control plane defined by GMPLS [1] and ASON [2]. The first choice provides fine-grane switching granularity but requires OEO conversion at intermediate IP routers, whereas the second choice comprises just the opposite.

A multi-layer hybrid architecture that efficiently combines the benefits of both switching paradigms is possible, but needs to be properly managed [3]. As noted in [4], new backbone designs require adaptive and agile mechanisms that properly manage the optical and electronic resources available. Note that this IP over WDM approach is evolutionary since both the IP routers and optical equipment are already deployed and maintained in the network, and new traffic demands are absorbed by the optical layer.

This work studies a Bayesian decisor algorithm that trades off two important metrics in the transmission of traffic demands: The end-to-end Quality of Service perceived and the relative operational cost of using the electronic and optical domains respectively. The general idea is that direct switching provides better e2e QoS but the optical resources are scarce to the network operator and must be used only when necessary, according to different criteria specified by the network administrator. In this light, the algorithm finds an optimal equilibria, whereby the optical (expensive) resources are used only when absolutely necessary, that is, only when otherwise (electronic switching) the e2e QoS degrades significantly reaching unacceptable levels. The algorithm is formulated following the Bayesian decision theory, thus providing a rigorous mathematical framework that can be adjusted by the network operator, yet keeping its original philosophy.

Concerning previous work on traffic grooming, the authors in [5] propose the minimisation of a single cost function defined in terms of wavelength occupation, but does not take into account the end-to-end QoS perceived by the flows. Our Bayesian decisor algorithm does provide a risk function that trades off two objectives: Cost and QoS. A similar formulation is found in [6], where an IP over WDM framework is defined but no deep analysis is performed on real network scenarios. Finally, in a previous work of ours [7], we propose a similar algorithm but its performance is only evaluated in a single-hop scenario, thus far from reality. This study aims at extending such previous formulation to a multiple node scenario and further evaluating its behaviour and dynamics.

The remainder of this work is organised as follows: Section II presents the multi-hop scenario whereby the Bayesian decision algorithm is to be applied. This section also introduces the risk function as it is defined in Bayesian decision theory, and computes its individual components: cost and utility (QoS perceived) to find the optimal decision. Section III shows the behaviour of the algorithm in such a multi-hop scenario and shows how to adjust the model parameters to trade off QoS and Cost. Finally, Section IV concludes this study and proposes future work.

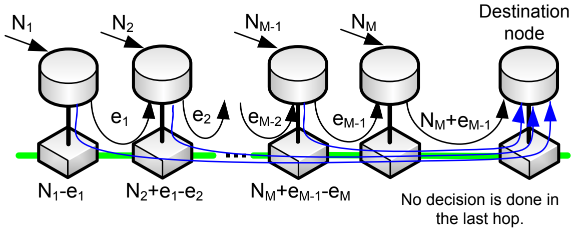

Fig. 1. Multi-hop scenario under study

II. PROBLEM STATEMENT: MULTI-HOP SCENARIO

Let us assume the multi-hop scenario proposed in Fig. 1, with $M+1$ multi-layer capable nodes. In this scenario, each node j (with $j = 1, \ldots, M$) is offered a number of N_j Label Switch Paths (LSPs) destined to the destination node in the Figure. Each node j decides the number e_j of LSPs that are electronically switched to the destination node, thus remaining o_j LSPs to be transmitted all-optically. The electronic transmission of LSPs implies traversing all intermediate nodes, thus suffering O/E and E/O conversions and electronic buffering (hence delay) at each of them. Optical switching implies the creation of an end-to-end (e2e) lightpath (from the source to the destination node) with no O/E conversion or delay experienced. Clearly, optical switching provides better QoS experienced by the LSPs (no delay at intermediate nodes), but requires extra resource consumption (the creation of new lightpaths per e2e optical switching). These two aspects of multi-layer switching (QoS perceived and extra Cost of using resources) must be traded off by the Bayesian decisor to find the optimum number of optically- and electronically-switched LSPs.

For notation purposes, let node 1 offer N_1 LSPs to the multi-layer Bayesian decisor. This decides to switch e_1 LSPs electronically (thus offered to node 2), and $o_1 = N_1 - e_1$ LSPs optically (Fig. 1). Node 2 therefore must decide the number e_2 of electronically switched LSPs among the total $N_2 + e_1$ offered, thus leaving $o_2 = N_2 + e_1 - e_2$ LSPs to go all-optically to the destination node. Following this reasoning, node j is offered $N_j + e_{j-1}$ total LSPs and, among them, e_j and $o_j = N_j + e_{j-1} - e_j$ are transmitted through the electronic (hop-by-hop switching) and optical domains (direct lightpath) respectively. Let us remark that an e2e ligthpath in the M-th node uses the same resources than a hop-by-hop service, hence the last hop makes no decision (thus $e_M = N_M + e_{M-1}$).

The next section defines the Risk function, which is based on the cost of using optical and electronic resources and the e2e delay experienced by the electronically switched LSPs.

A. Risk function definition

Let $\vec{e} = \{e_1, e_2, \ldots, e_{M-1}\}$ denote the decision vector which gives the number of LSPs transmitted over the IP (electronic) layer at each hop j. The queueing delay experienced by electronically-switched LSPs or flows at hop j (that is, delay from node j to node $j+1$) x_j depends on the load offered e_j. Thus, the e2e delay experienced by a given LSP offered at node j is $x_j^{e2e} = \sum_{i=j}^{M} x_i$ since it must traverse the subsequent nodes until destination (with their respective delays). With these parameters, we define the following Bayesian Risk function [7], [8]:

$$\begin{aligned} R(\vec{e}, x_j^{e2e}) &= K_c C_T(\vec{e}) - K_u \sum_{j=1}^{M} \mathbb{E}_x \left[U(x_j^{e2e}) \right], \\ & x_j^{e2e} \geq 0 \end{aligned} \quad (1)$$

where $C_T(\vec{e})$ and $U(x_j^{e2e})$ refer to the utilisation cost and utility function associated to the decision vector $\vec{e}$. Both quantities are weighted by constants K_c and K_u. Clearly, the goal is to find the optimum decision vector $\vec{e}^*$ that minimises the Bayesian Risk given by Eq. 1.

Since the decision vector $\vec{e}^*$ gives the number of LSPs switched in each domain (optical and electronic), the Bayesian Risk finds a trade-off between the "Utility" perceived by the traffic sent through the electronic domain and the cost related with the utilisation of the optical and electronic resources. The following explains how to compute $C_T(\vec{e})$ and $\mathbb{E}_x \left[U(x_j^{e2e}) \right]$.

B. Cost of using resources

This function accounts the cost C_e of using hop-by-hop connections (electronic switching) and the cost C_o of using end-to-end ligthpaths (optical switching) which, for a decision vector $\vec{e}$, is given by:

$$C_T(\vec{e}) = C_e(\vec{e}) + R_{\text{cost}} C_o(\vec{e}) \quad (2)$$

where R_{cost} is the relative cost of using the optical and electronic resources. In other words, an optical lightpath is R_{cost} times more expensive than the same connection in the electronic layer. Note that R_{cost} is not a monetary cost but a metric that helps network operators decide how valuable their optical resources are with respect to the already deployed IP layer.

Cost computation has been chosen to follow the next design premises: (1) LSPs should be switched in the electronic domain while their utility perceived is correct, hence the cost of using electronic resources is cheaper than that of using optical resources, for the same amount of traffic (routers are already deployed); and (2) if an optical bypass is to be set up, the longer it is, the better (less cost), that is, the cost of long connections should be lower than short optical by-pass connections. Thanks to this cost model, only the necessary e2e optical connections are created, and this occurs when the IP layer do not provide the necessary utility to the traffic.

Following these premises, we define the cost of transmitting an LSP optically per hop as $\frac{k+1}{k}$, where k is the length of the optical by-pass (that is, a lightpath created from node j to the destination node is of length $k = M + 1 - j$). Note that this series is strictly decreasing since $\frac{k+1}{k} > \frac{l+1}{l}$, $\forall k < l$, giving a cheaper cost per hop the longer the lightpath is, thus promoting the creation of long e2e by-pass optical connections in the network. It is worth noticing that, in the scenario proposed in Fig. 1, the longest (and cheapest) lightpath possible is of cost $\frac{M+1}{M}$, and it is the cheapest one since $\frac{M+1}{M} < \frac{k+1}{k}$, $\forall k < M$.

In conclusion, the optical cost of sending i LSPs through k hops is $\frac{k+1}{k} \times k \times i = (k+1) \times i$. The definition of the electronic and optical cost functions are as follows:

$$\begin{aligned} C_e(\vec{e}) &= \sum_{j=1}^{M} 2e_j \quad (3) \\ C_o(\vec{e}) &= ((M+1)(N_1 - e_1) \\ &\quad + \sum_{j=2}^{M-1} (M-j+2)(N_j - e_j + e_{j-1}) \Bigg) \quad (4) \end{aligned}$$

According to the previous definitions, optical resources are R_{cost} times more expensive than electronic in the case of one-hop switching. The one-hop electronic cost has been assumed of value 2 ($k = 1$), and $\frac{(M+1)}{M}$ is the cheapest optical cost per hop in this scenario. According to this, R_{cost} must satisfy $R_{\text{cost}} > \frac{2 \cdot M}{M+1}$ to ensure that the cheapest optical lightpath is more expensive than its electronic counterpart.

C. Queueing delay model

The utility functions in Eq. 1 are defined based on delay as QoS metric. For this reason, it is necessary to define a queueing delay model. As shown in previous studies, the queueing delay of routers fed by self-similar traffic can be accurately characterised by a Weibull distribution, see [9], [10], [11], For a single queue, such a probabilistic delay is a function of the network load (number of electronically switched LSPs), lightpath capacity C and Hurst coefficient H, as noted in [9]:

$$\begin{aligned} p(x) &= \frac{s}{r^s} x^{s-1} \exp\{-\left(\frac{x}{r}\right)^s\}, x \geq 0 \quad (5) \\ s &= 2 - 2H \\ r &= \frac{1}{C}\left(\frac{2K(H)^2 ame}{(C-me)^{2H}}\right)^{\frac{1}{2-2H}} \end{aligned}$$

where m is the average input traffic per LSP and a is a variance coefficient such that $am = \sigma^2$ (with σ^2 being the input traffic variance). Clearly, the more electronic LSPs e, the more delay experienced in the queue.

D. Utility functions definition

The utility function applied to a decision vector $\vec{e}$ gives a metric for the delay experienced by the electronically-switched LSPs, such that, the more delay experienced by them, the less utility achieved. The electronically-switched LSPs are assumed to experience some degree of delay, since they must traverse several hops with their respective electronic queues. On the other hand, the delay experienced by the optically-switched LSPs is assumed negligible compared to the electronic delay, since optical LSPs are provided a dedicated e2e path. Such an electronic delay is calculated based on the load level of a queue fed with self-similar traffic, as explained in Section II-C. Once the e2e electronic delay is obtained, the utility function operates to derive a utility metric following one of these Class of Service (CoS) utility models [7]: average delay, hard real-time and elastic utilities, as follows:

1) Average delay-based utility (U_{mean}): This utility is defined as: $U_{mean}(x_j^{e2e}) = -x_j^{e2e}$ which, after applying the expectation operator $\mathbb{E}_x$ of Eq. 1, provides a utility function based on the average e2e delay experienced by the electronically-switched LSPs. This value is computed as:

$$\begin{aligned} \mathbb{E}_x[U_{\text{mean}}(x_k^{e2e})] &= \mathbb{E}_x[-x_k^{e2e}] = -\sum_{j=k}^{M} \mathbb{E}_x[x_j] \\ &= -\sum_{j=k}^{M} \left\{ r\Gamma\left(1 + \frac{1}{s}\right)\right\} \quad (6) \end{aligned}$$

This utility function can be used for best-effort services, whereby great service interactivity provides high utility values (sending an email instantly), but this utility function does not excessively penalyse if such interactivity is low. The following utility functions are more restrictive with delay.

2) Hard real-time utility (U_{step}): Some applications tolerate very well a certain e2e delay value until a given delay threshold, but are completely useless if such a threshold is exceeded. Examples of these are: on-line gaming, back-up services and grid applications. The ITU-T recommendation Y.1541 [12] and the 3GPP recommendation S.R0035 [13] define Classes of Service (CoS) based on such thresholds. According to this, we propose the following step utility function to deal with such scenarios:

$$U_{\text{step}}(x_j^{e2e}) = \begin{cases} 1, & \text{if } x_j^{e2e} < T_{\max} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where the threshold $T_{\max}$ depends on the service or application.

After applying the expectation operator $\mathbb{E}_x$ to U_{step}, it yields [7]:

$$\mathbb{E}_x[U_{\text{step}}(x_k^{e2e})] = \mathbb{E}_x[U_{\text{step}}(\sum_{j=k}^{M} x_j)] = P(x_k^{e2e} < T_{\max}) \quad (8)$$

The calculation of the e2e delay expectation requires the convolution of the queueing delay pdf, which it is not possible to obtain analytically. However, we can approximate the e2e delay (x^{e2e}) by a Gaussian distribution, assuming that the per-hop delays are independent. The moments of such a Gaussian pdf are computed by the Weibull delay assumption (Eq. 5):

$$P(x_j^{e2e} < T_{\max}) \sim N(\sum_{i=j}^{M} \mu_i, \sqrt{\sum_{i=j}^{M} \sigma_i^2}) = N(\mu_j^{e2e}, \sigma_j^{e2e}) \quad (9)$$

3) Elastic utility (U_{exp}): Other applications, such as voice transmission, experience slow service degradation with increasing delay, until a threshold delay is reached, following the "E model" [14]. We propose the following elastic utility function to deal with delay-sensitive applications:

$$U_{\exp}(x_j^{e2e}) = \lambda e^{-\lambda x_j^{e2e}}, \quad x_j^{e2e} \geq 0 \quad (10)$$

where λ refers to decay ratio of the exponential function. This utility function lies somewhere in between the previous two, whereby excessive delays are highly penalised, but not that much as in the hard real-time utility case.

Finally, the value of λ is chosen such that $\alpha = 50\%$ of the total utility lies before a delay threshold T_{max}:

$$\lambda = \frac{1}{T_{\text{max}} \log(1-\alpha)} \tag{11}$$

this value can be obviously adjusted for a given α.

Finally, after assuming the Gaussian approximation of Eq. 9 for computing e2e delays, the expected utility obtained in this case is given by:

$$\begin{aligned} \mathbb{E}_x[U_{\text{exp}}(x_j^{e2e})] &= \mathbb{E}_x[\lambda e^{-\lambda x_j^{e2e}}] \qquad (12) \\ &= \int_{-\infty}^{\infty} \lambda e^{-\lambda x_j^{e2e}} N(\mu_j^{e2e}, \sigma_j^{e2e}) dx_j^{e2e} \end{aligned}$$

This integral can be solved completing the square, achieving the following expression:

$$\mathbb{E}_{x_{e2e}}[U_{\text{exp}}(x_j^{e2e})] = \lambda e^{\frac{\sigma_j^{2\ e2e}\lambda^2 - 2\mu_j^{e2e}\lambda}{2}} \tag{13}$$

III. NUMERICAL RESULTS AND DISCUSSION

Next, the Bayesian decisor is evaluated in the scenario depicted in Fig. 1 with $M = 3$ hops. We assume the following parameter values: 2.5 Gbps of lightpath capacity fed with standard VC-3 LSPs of 34.358 Mbps bitrate each. The number of incoming flows in the last node is $N_3 = 0$. The bandwidth standard deviation is chosen such that $\frac{\sigma}{m} = 30\%$ and the Hurst parameter selected is $H = 0.6$ (according to [15]). The T_{max} value is set to $80ms$ for U_{exp} and $T_{\text{max}} = 5ms$ for U_{step}, since the QoS restrictions are more stringent in the latter case. The value of $R_{\text{cost}} = 2$ by default. K_c and K_u are constants that define the decision when the system operates at maximum network load (that is, $N_{max} = \lfloor C/m \rfloor$). Thanks to these constants it is possible to fix the occupation of the optical and electrical link in the worst case. In our numerical experiments, for $N_{max} = \lfloor C/m \rfloor = 72$ incoming LSPs, the hop-by-hop electronic connection transmits 70% of the traffic, that is 50 LSPs. This policy can be adjusted by the network operator as necessary.

A. *Decisor dynamics experiment*

The level curves of the Risk function help us to see how the function changes with the incoming traffic. Fig. 2 (left) shows the U_{mean} level curves for $N_1 = 72$ and no cross traffic at node 2 ($N_2 = 0$). Since this is the normalisation working point, the algorithm decides to send $N_{omax} = 50$ LSPs through the IP layer. Fig. 2 (right) illustrates the decision when node 2 injects some cross traffic, more specifically $N_2 = 10$ LSPs. In this situation, the decisor changes its behaviour by sending $o_1 = 35$ LSPs through the optical layer from node 1 to the destination node, which gives $e_1 = 72 - 35 = 37$ LSPs through the electronic domain. These 37 LSPs are added to the $N_2 = 10$ offered at node 2, which are transmitted electronically to the destination node.

It is worth noting that, since node 2 is closer to the destination node than node 1, its QoS restrictions are more permissive than if the same amount of traffic was offered at node 1.

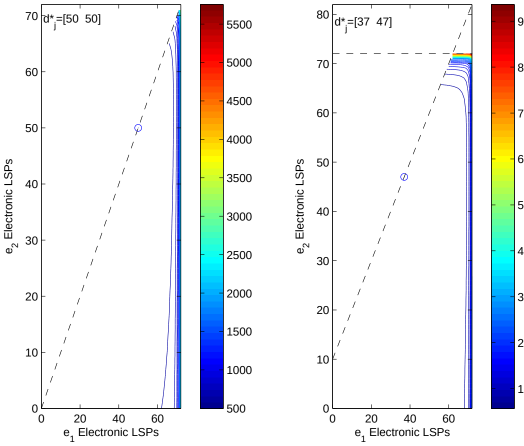

Fig. 2. Level curves examples (U_{mean})

The level curves of the Risk function for U_{exp} and U_{step} are not included for brevity purposes, but the next experiments examine the behaviour of the decision algorithm for such utility functions.

1) Traffic increment in the first node without cross-traffic: The next experiment shows how the algorithm changes its decision when the first node increases the amount of LSPs offered to the system and there is no cross-traffic ($N_2 = 0$). Fig. 3 (left) shows the number of flows sent through the electronic and optical domains at each hop, for the average delay utility case (U_{mean}). As shown, all traffic flows are sent through the IP layer until the utility given to the flows is smaller than the cost of establishing a new e2e connection, which occurs when $N_1 \geq N_{omax}$. At this point, a direct lightpath (first lightpath) from the first node to the destination node is created, as shown in the figure. After this, the network load keeps increasing (more LSPs offered at node 1) and, after some time (when the delay experienced at the second hop is excessive), a second lightpath at node 2 is created for incoming LSPs.

Fig. 3 (center and right) illustrates the same experiment but using the U_{step} and U_{exp} utility functions instead. The U_{exp} utility function behaves very similarly to U_{mean}. However, when the U_{step} is employed, the system forces all electronically-switched LSPs in the second node to be switched over the second lightpath, once this is created. The step utility is shown to be more QoS aware than both the exp and mean utility functions.

2) First node constant rate and second node load increment: This experiment evaluates the decision when the load offered to the first node is constant ($N_1 = 10$) and the second node sends a variable number N_2 of LSPs. Fig. 4 depicts the amount of traffic sent using the electronic and optical layers in both hops. When the second node gets saturated (QoS

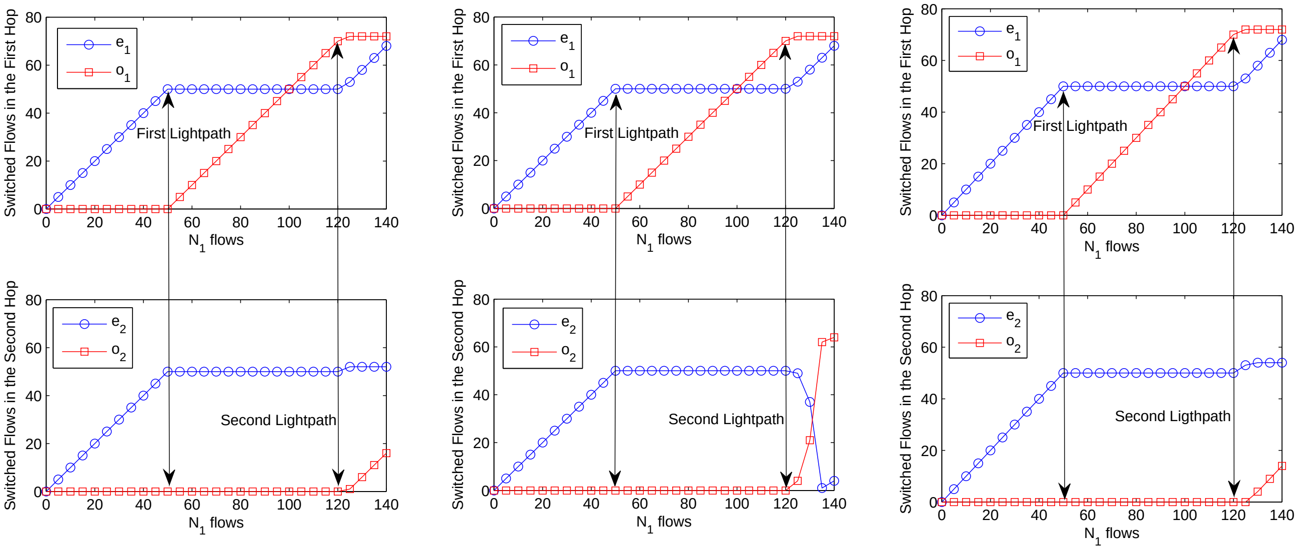

(a) Umean (b) Ustep (c) Uexp

Fig. 3. LSPs sent through the electronic and optical layers at nodes 1 and 2 when the load in the first node increases

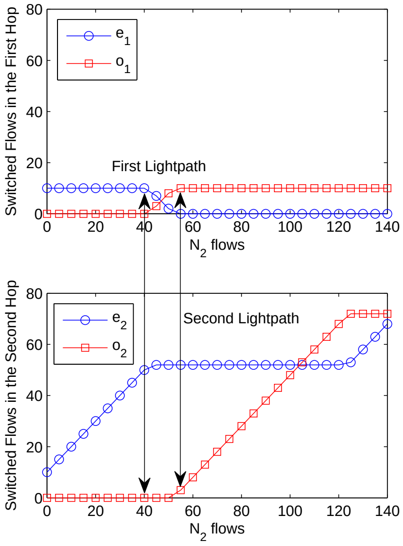

Fig. 4. LSPs sent through the electronic and optical layer when the load in the second node increases (U_{mean} only)

degraded), the first node decides to send its 10 LSPs using a direct e2e lightpath (first lightpath). It is worth noticing that a lightpath is created at the first node, rather than at the second node. This behaviour occurs thanks to the cost function, which favors the creation of long ligthpaths. However, since the traffic offered at node two N_2 keeps increasing, the bayesian decisor establishes a second e2e ligthpath at the second node.

B. On the influence of the decisor's parameters

1) Influence of the utilisation cost (R_{cost}): Table I shows the optimal decision for different values of R_{cost} and the three utility functions. Remark that R_{cost} satisfies the condition: $R_{\text{cost}} > \frac{2*M}{M+1} = \frac{3}{2}$ to make the cheapest (longest) lightpath more expensive than its electronic counterpart. The results show that the more expensive the optical resources are (large values of R_{cost}), the fewer LSPs are routed using the optical domain, as expected. When U_{mean} and U_{exp} are used, the value of R_{cost} indeed decides the number of LSPs switched through each domain. For example, with U_{exp}, $e_1^* = 32$ LSPs are switched over the electronic layer for $R_{\text{cost}} = 1.6$, while for $R_{\text{cost}} = 3$, we have $e_1^* = 58$. On the other hand, the results obtained for the U_{step} function are different than for U_{mean} and U_{exp}. In this case, the decision does not vary significantly with respect to R_{cost} (Table I), since the decision is mostly determined by the QoS parameters.

		$N_1 = 60, N_2 = 0$			$N_1 = 60, N_2 = 10$		
		U_{mean}	U_{step}	U_{exp}	U_{mean}	U_{step}	U_{exp}
$R_{\text{cost}} = 1.6$	e_1^*	33	50	32	17	41	16
	e_2^*	33	50	32	27	51	26
$R_{\text{cost}} = 2$	e_1^*	50	50	50	37	42	37
	e_2^*	50	50	50	47	52	47
$R_{\text{cost}} = 3$	e_1^*	58	51	58	54	49	54
	e_2^*	58	51	58	55	52	55

TABLE I
OPTIMAL DECISIONS WITH THE VARIATION OF THE R_{COST} PARAMETER

2) Study of delay QoS threshold (T_{max}): This section presents the decision results for changing T_{max} for offered traffic $N_1 = 60$ and $N_2 = 10$ fixed. As previously stated in Section II-A, the QoS parameter (T_{max}) is only introduced for the elastic (U_{exp}) and hard-real time (U_{step}) utility functions. Therefore, U_{mean} is not studied in this section.

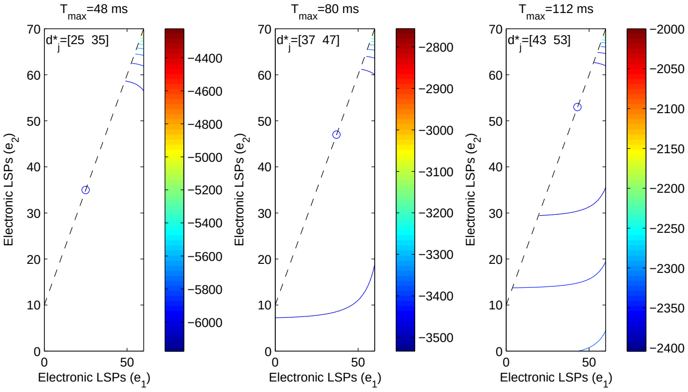

Fig. 5. Variation of the $T_{\max}$ parameter (U_{exp})

Fig. 5 illustrates the optimal decision for U_{exp} when $T_{\max}$ varies from $48ms$ to $112ms$, which is a 40% of variation from $80ms$. In light of the results, we observe that the network operator can tune the number of LSPs to be sent through the optical layer by changing $T_{\max}$ value. If flows are subject to coarser QoS constraints, the Bayesian decisor sends more LSPs over the electronic layer.

For the U_{step} function, the results are the following: for $T_{\max} = 3ms$ is $\vec{e} = \{37, 47\}$, for $T_{\max} = 5ms$ is $\vec{e} = \{42, 52\}$ and for $T_{\max} = 7ms$ is $\vec{e} = \{45, 55\}$. The variation from $3ms$ to $7ms$ is a 40% from $5ms$ to make a fair comparison. Table I showed that U_{step} is not very sensible to R_{cost}, but it is to $T_{\max}$ (the QoS parameter). The reason is that this parameter is related to the e2e QoS performance experienced by the LSPs.

IV. SUMMARY AND FUTURE WORK

This work proposes a novel methodology to deal with the utilisation of the electronic and optical layers in a multi-hop scenario with multi-layer capable routers. Essentially, this includes the definition of a multi-hop Bayesian decisor which decides the amount of traffic routed through the optical and electronic domains, and its behaviour is explained. Thanks to the $T_{\max}$ and R_{cost} parameters, the network operator is provided with a means to define QoS and Relative optical/electronic cost aware metrics, which can be further applied to define traffic engineering mechanisms.

As future work, this decisor will be studied in a full network topology, on attempts to define a full risk-oriented routing mechanism. Finally, the provisioning of multiple services in the same network scenario is an open issue of interest to be studied in the future.

ACKNOWLEDGEMENTS

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.

The authors would also like to acknowledge the support of the Spanish projects: "Multilayer Networks: IP over Transport Networks" (TEC2008-02552-E) and DIOR (TEC-2006-03246).

REFERENCES

[1] E. Mannie and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) control," Internet Engineering Task Force, Tech. Rep., 2004.

[2] *Architecture for the Automatically Switched Optical Network (ASON)*, ITU-T Rec. G.8080/Y.1304, November 2001.

[3] K. Sato, N. Yamanaka, Y. Takigawa, M. Koga, S. Okamoto, K. Shiomoto, E. Oki, and W. Imajuku, "GMPLS-based photonic multilayer router (Hikari router) architecture: an overview of traffic engineering and signaling technology," *IEEE Communications Magazine*, vol. 40, no. 3, pp. 96–101, mar 2002.

[4] C. M. Gauger, P. J. Kuhn, E. Van Breusegem, M. Pickavet, and P. Demeester, "Hybrid optical network architectures: Bringing packets and circuits together," *IEEE Comms. Magazine*, vol. 44, no. 8, pp. 36–42, Aug. 2006.

[5] B. Puype, Q. Yan, S. Colle, D. De Maesschalck, I. Lievens, M. Pickavet, and P. Demeester, "Multi-layer traffic engineering in data-centric optical networks," in *Proceedings of Optical Networking Design and Modeling (ONDM)*, 2003.

[6] A. Elwalid, D. Mitra, and Q. Wang, "Distributed nonlinear integer optimization for data-optical internetworking," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1502–1513, Aug. 2006.

[7] V. López, J. A. Hernández, J. Aracil, J. P. F. Palacios, and O. G. de Dios, "A Bayesian Decision Theory Approach for the Techno-Economic Analysis of an All-Optical Router (extended version)," *Computer Networks*, vol. 52, no. 10, pp. 1916–1926, Jul 2008.

[8] S. French and D. Ríos Insúa, *Statistical decision theory.* Oxford University Press Inc., 2000.

[9] I. Norros, "On the use of fractional Brownian motion in the theory of connectionless networks," *IEEE J. Selected Areas in Communications*, vol. 13, no. 6, pp. 953–962, Aug 1995.

[10] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, and C. Diot, "Measurement and analysis of single-hop delay on an ip backbone network," *IEEE J. Selected Areas in Communications*, vol. 21, no. 6, pp. 908–921, Aug 2003.

[11] J. A. Hernández and I. W. Phillips, "Weibull mixture model to characterise end-to-end Internet delay at coarse time-scales," *IEE Proc. Communications*, vol. 153, no. 2, April 2005.

[12] ITU-T, "ITU-T Recommendation Y.1541 - Network Performance Objectives for IP-Based Services," Feb. 2003.

[13] 3GPP, "3GPP Recommendation S.R0035-0 v1.0. - Quality of Service," Sep. 2002.

[14] ITU-T, "ITU-T Recommendation G.107 : The E-model, a computational model for use in transmission plannings," Mar. 2005.

[15] R. G. Clegg, "Markov-modulated on/off processes for long-range dependent Internet traffic," *ArXiv Computer Science e-prints*, Oct. 2006.