Interdomain RWA based on stochastic estimation methods and adaptive filtering for optical networks

M. Yannuzzi, X. Masip-Bruin, S. Sánchez-López, E. Marin Tordera, J. Solé-Pareta, and J. Domingo-Pascual
Advanced Network Architectures Lab (CRAAX/CCABA), Technical University of Catalunya, Barcelona, Spain
{yannuzzi, xmasip, sergio, eva, pareta, jordid}@ac.upc.edu

***Abstract*- This paper presents a RWA strategy based on the stochastic estimation of the Effective Number of Available Wavelengths (ENAW) along interdomain paths. We propose an approximate model to roughly estimate the ENAW on the paths across multiple domains, and then refine this estimation by means of observations and an adaptive prediction-correction Kalman filtering process. Unlike traditional Kalman filters, which use noisy measurements as their observations, we use the information contained in the routing updates as noisy measurements of the existent wavelength occupancy along the paths. Based on these observations, we estimate the ENAW on the candidate paths, and use this information to influence the RWA decision when the number of available wavelengths makes a lightpath request prone to be blocked. The approximate model we are proposing in this paper is based on a noisy extension of a simplified model derived for two domains. Our results validate the usefulness of the model, and confirm that by estimating the wavelength occupancy prior to the RWA decision the blocking ratio can be considerably reduced.**

I. Introduction

A key problem in any Routing and Wavelength Assignment (RWA) strategy for multi-domain wavelength-routed optical networks is how to assess the wavelength availability on a candidate interdomain lightpath. The reason for this is threefold. In the first place, routing protocols devised for large-scale multi-domain networks, require the aggregation of network management information at domain boundaries so as to keep the protocol scalable. Second, for confidentiality reasons, each domain discloses only the amount of information matching its own requirements and business policies. And third, keeping multi-domain networks fully synchronized with perfectly updated information about the RWA aggregates, seems, at least at present, also unfeasible for scalability reasons [1]. These three factors contribute to make the wavelength availability information on interdomain paths quite "noisy".

In the current Internet, BGP routers only advertise highly aggregated *reachability information*, turning each downstream domain along an interdomain path into a black-box. This approach has proven to be highly scalable, but unfortunately it imposes several strong limitations, especially, in terms of end-to-end QoS and interdomain Traffic Engineering (TE) [2]. Accordingly, it is widely accepted now that in addition to reachability information, neighboring domains should become capable of exchanging some aggregated *path state information*.

Future on-demand lightpath provisioning networks will natively demand TE, QoS, and Quality of Resilience (QoR) capabilities between different domains. Given that it is quite unlikely that a domain blindly trust the decisions made by downstream domains, simply managing reachability information is not enough from a network operator standpoint. This has leveraged the proposals of different network state aggregation schemes and updating policies at the interdomain level for wavelength-routed optical networks [1, 3]. Thus, in this work we consider that neighboring domains are able to exchange both reachability, and partial path state information consisting of aggregated wavelength availability and aggregated load information. This approach is aligned with the main ideas in [1], but unlike that work, we do not focus on the aggregation scheme. Our focus is on the RWA algorithm, and how to influence its decision using forecasted estimates of the wavelength availability contained in "noisy" information aggregates.

For the information exchange we use the Inter-Domain Routing Agents (IDRA) that we proposed in [3]. Each domain or Autonomous System (AS) allocates one or more of these agents, which are the ones in charge of computing paths, and exchanging routing updates containing reachability and path state information. The Fig. 1 depicts the approach. On the one hand, the Network Reachability Information (NRI) messages are triggered each time a new route is available or a known one becomes unavailable. On the other hand, aggregated Path State Information (PSI) associated with the destinations contained in the NRI messages is exchanged between the IDRAs. To this end, we take advantage of the Keepalive messages exchanged between the IDRAs, and extend them with the purpose of conveying PSI. This allows an IDRA to send a non-dummy Keepalive message when relevant PSI needs to be updated.

In this framework, we propose a linear stochastic model to roughly estimate the number of available wavelengths along an interdomain path between two routing updates, and then tune this coarse-grained estimation by means of an adaptive and discrete-time prediction-correction Kalman filter. We use this estimation to influence the IDRA's RWA decision when the number of available wavelengths on *all* candidate paths makes a lightpath request prone to be blocked. Our results show that this estimation considerably improves the performance of the RWA algorithm in terms of the lightpath blocking ratio. It is worth highlighting that Kalman filters have been widely used in different disciplines, like adaptive control, ATM, and recently in IP/MPLS networks, given their optimal estimation-prediction error characteristics [4-6].

This work was partially funded by the Spanish Ministry of Education under contract TEC2005-08051-C03-01, the Catalan Research Council under contract 2005-SGR00481, and the European Commission through ePhoton/One+ under contract FP6-027497.

The rest of the paper is structured as follows. In Section II we describe the NRI and PSI exchanged between the IDRAs. In Section III we introduce the stochastic model. In Section IV we present the Kalman-based predictor-corrector we use in order to adaptively tune the outcome of the previous model, and the way in which it is utilized to influence the RWA algorithm. Our simulation results are shown in Section V. Finally, Section VI concludes the paper.

II. Reachability and Aggregated Path State Information Exchange

The IDRAs are the ones responsible for carrying interdomain routing information, and deciding within each AS which path is the best to reach any known destination[1]. Thus, the RWA algorithm that we propose in this paper runs in the IDRAs. In this particular work we assume that the Optical Cross Connects (OXCs) do not perform wavelength conversion, so each lightpath computed by an IDRA is subject to the wavelength continuity constraint. The rest of this section describes the NRI and the aggregated PSI conveyed by the IDRAs in Fig. 1.

A. Network Reachability Information (NRI) exchange

Let L, F, and Ω denote the number of links, the number of fibers per-link, and the number of wavelengths (colors) per-fiber respectively, at each destination OXC within an AS. For the sake of simplicity we assume that all destination OXCs are identical, and that each network sinking traffic consists of only one destination OXC. Thus, $LF\Omega$ is an upper bound of the number of available wavelengths to reach any given destination. Each AS may select (according to its local TE and routing policies) the particular subset of wavelengths that can be used by an upstream domain to reach the local destination networks. Consequently, the reachability information contained in the NRI messages consists of destination networks, and a set of pairs $R=\{(\Lambda_1, M(\Lambda_1)),\ldots,(\Lambda_N, M(\Lambda_N))\}$, wherein Λ_i denotes a particular wavelength (color), and $M(\Lambda_i)$ denotes the maximum multiplicity of Λ_i. Clearly, $N \leq \Omega$, and $M(\Lambda_i) \leq LF\ \forall\ i$. For each destination network, a transit AS may filter and advertise a subset of R to its upstream domains, or simply retransmit the NRI messages received. When a new destination network becomes available, or an already known one becomes unavailable, the NRI messages are triggered immediately by an IDRA. In any other case, the NRI should only change over rather large timescales compared to the PSI, according to the local optimizations and TE actions performed by the different domains.

B. Path State Information (PSI) exchange

The PSI is composed by *aggregated wavelength availability* and *aggregated load information*. Each IDRA advertises PSI messages by aggregating the following three pieces of information: i) the intradomain PSI; ii) the PSI related to the interdomain links toward its downstream domains; iii) the already aggregated PSI contained in the interdomain advertisements received from downstream domains. In the sequel we will describe how this aggregation process is done.

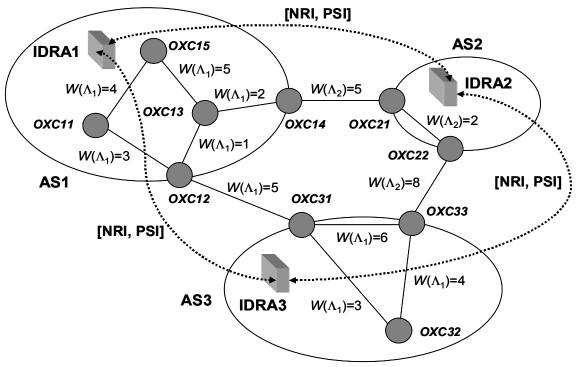

Fig.1. Network Reachability Information (NRI) and Path State Information (PSI) exchange between the IDRAs. For scalability reasons, the agents are decoupled from the Optical Cross Connects (OXCs).

Aggregated Wavelength Availability Information: Let r and q be a pair of OXCs inside an AS, $P(r, q)$ be a candidate path between r and q, and l be a link within the path $P(r, q)$. An IDRA computes the Effective Number of Available Wavelengths (ENAW) of type Λ_i between the OXCs r and q as follows:

$$W_{r,q}(\Lambda_i) = \max_{P(r,q)}\left\{\min_{l\in P(r,q)}\left[W_l(\Lambda_i)\right]\right\} \tag{1}$$

The rationale in (1) can be easily interpreted by means of the Fig. 1. For instance, in AS1 the ENAW of type Λ_1 between the nodes *OXC15* and *OXC12* is $W_{15,12}(\Lambda_1)=3$. This is because from the two possible paths between these nodes, the path that goes through *OXC13* has a minimum $W_{15,12}(\Lambda_1)=1$, whereas the one that goes through *OXC11* has a minimum $W_{15,12}(\Lambda_1)=3$. Then, the maximum between both of them is 3. The ENAW given in (1) is especially important between two border OXCs in a transit domain, since it captures the practical availability of the wavelength Λ_i within the domain. In addition, (1) offers highly aggregated network state information, so this is the intradomain portion of the wavelength availability component of a PSI aggregate.

For the interdomain portion, each IDRA knows which wavelengths are actually being used on its interdomain links, and it also knows which wavelengths are effectively available downstream through the PSI advertisements received from neighboring IDRAs. Let $W_{lb,rb}(\Lambda_i)$ denote the number of available wavelengths of type Λ_i in the interdomain link between the local border node lb, and a remote border node rb. For instance, in Fig. 1 the IDRA1 in AS1 is aware that $W_{12,31}(\Lambda_1)=5$. Similarly, let $W^{adv}_{rb,d}(\Lambda_i)$ denote the ENAW of type Λ_i between the remote border node rb and the destination node d, advertised by the downstream IDRA in rb's domain. Using these two interdomain components and (1), an IDRA advertises upstream that the ENAW between a local border node lb and a distant destination node d is:

[1] Signaling issues are left out of the scope of this paper.

$$W_{lb,d}^{adv}(\Lambda_i) = \min\left\{W_{lb,lb'}(\Lambda_i), W_{lb',rb}(\Lambda_i), W_{rb,d}^{adv}(\Lambda_i)\right\} \quad (2)$$

For instance, in Fig. 1 the IDRA1 advertises to the IDRA2 that the ENAW of type Λ_1 to reach *OXC32* is $W_{14,32}^{adv}(\Lambda_1) = \min\left\{W_{14,12}(\Lambda_1), W_{12,31}(\Lambda_1), W_{31,32}^{adv}(\Lambda_1)\right\} = \min\{2,5,4\} = 2$.

Aggregated Load Information: This comprises two sets of state information, namely, aggregated costs and aggregated blocking ratios. On the one hand, an additive cost is associated with each candidate (path, wavelength) pair. This cost reflects the current load in the availability of wavelengths in a path, allowing an IDRA to tiebreak when two or more candidate paths offer almost the same ENAW. The cost associated with a candidate path $P(s, d)$ between a local node s and a distant node d for wavelength type Λ_i is computed by an IDRA as follows:

$$C_{P(s,d)}(\Lambda_i) = \begin{cases} H\left[\dfrac{1}{\min\left[W_{s,lb'}(\Lambda_i), M(\Lambda_i)\right]} + \dfrac{1}{\min\left[W_{lb',rb}(\Lambda_i), M(\Lambda_i)\right]}\right] + \dfrac{C_{P(rb,d)}^{adv}(\Lambda_i)}{H^{adv}} \\ \infty \quad \text{if } W_{s,lb'}(\Lambda_i) = 0 \text{ or } W_{lb',rb}(\Lambda_i) = 0 \end{cases} \quad (3)$$

Wherein, H is the number of hops from s to d considering each intradomain sub-path as just one hop. Similarly, H^{adv} is the number of hops between the remote border node rb and the destination node d, advertised by the downstream IDRA in rb's domain. The term $C^{adv}{}_{P(rb,d)}(\Lambda_i)$ denotes the cost from rb to d advertised by the downstream IDRA. The ∞ in (3) reflects the lack of resources to handle a connection between the nodes s and d for a particular wavelength Λ_i. If this is the case, an IDRA will remove from the NRI field of its advertisements all the destinations that were reachable through the path $P(s, d)$ for Λ_i.

The rationale in (3) is that the cost increases when the ENAW along an interdomain path decreases. Likewise, the cost increases when the length of an interdomain path increases, so an IDRA will "generally" choose the ($P(s, d)$, Λ_i) pair with the lowest cost.[2] It is worth highlighting that different candidate paths offering the same ENAW will frequently have different costs (loads). For instance, in Fig. 1 *OXC14* can reach *OXC33* both through AS1 and through AS2. The ENAW of type Λ_1 through AS1 is $W_{14,33}(\Lambda_1)$=2, and this is also the case for Λ_2 through AS2, i.e., $W_{14,33}(\Lambda_2)$=2. From (3), it can be easily shown that from these two paths, the IDRA1 prefers the one through *OXC21* given that Λ_2 is less loaded (notice that H=3 for both candidate paths).

The second type of load information contained in a PSI message is an ordered sequence of aggregated Blocking Ratios (BRs) coming from downstream domains. Our approach is that each domain j appends in the BR sequence its $BR_j(d)$, which corresponds to the ratio of path requests toward a destination d that have been blocked due to the lack of resources inside domain j. Each domain computes and updates its local BRs on a reasonable time-basis (in the order of several minutes, hourly, daily, etc), so that (1 - $BR_j(d)$) approximately represents the probability of traversing domain j while trying to reach destination d.[3] In realistic settings it is expected that each BR in the path sequence remains low and its variations shouldn't be significant. We anticipate that the BR sequence will aid in development of the model in Section III.

In sum, the path state information received by an IDRA is composed as follows:

$$PSI(d) = \left[\left\{W_{rb,d}^{adv}, \left(C_{P(rb,d)}^{adv}, H^{adv}\right)\right\}_{\Lambda_i}, \left\{BR_j(d)\right\}\right] \quad (4)$$

In Section V we show that when the RWA algorithm exploits the highly aggregated load information in (3), it is possible to obtain significant benefits in terms of the number of blocked interdomain calls. However, when the ENAW along *all* candidate paths is low, cost-based RWA algorithms yield blockings which can be considerably improved. The reason for this is that the load information given in (3) does not take into account the traffic demands. This is precisely what we address in the next two sections.

III. Stochastic Estimation of the Wavelength Availability Along Interdomain Paths

In state dependent circuit-switched networks the occupancy and the traffic arrival rates are typically coupled to each other, since the occupancy determines the traffic carried by the network and the carried traffic determines in turn the occupancy [7]. As a consequence, models developed to obtain explicit forms of the occupancy are highly complex, and typically involve multidimensional Markov processes leading to a set of coupled non-linear equations [8]. Unfortunately, the occupancy does not have a close-form expression, so numerical evaluations and complex iterations are needed in order to find either the blocking probabilities or the occupancy along the paths.

In this paper we propose a quite different approach. We aim at relaxing the model complexity by deriving a noisy *linear* model to roughly estimate the availability of wavelengths along the paths, and rely on a Kalman-based predictor-corrector to refine the previous estimation. Our model is based on a noisy extension of the estimation obtained for the two domains in Fig. 2, so we will first focus on this case.

The Fig. 2 shows a source domain AS1 and a directly connected destination domain AS2 consisting of a single OXC *OXC2*. The interdomain calls from AS1 to AS2 for wavelength Λ_i are assumed to be Poisson with exponentially distributed arrival rate λ. The duration of these calls are also assumed to be exponentially distributed with departure rate μ. For the sce-

[2] The term "generally" reflects the fact that when the ENAW is low, the path selection will be driven by the Kalman estimation process rather than by cost.

[3] We consider that both network operators and customers can benefit from this approach, which is aligned with some of the main ideas proposed in [1].

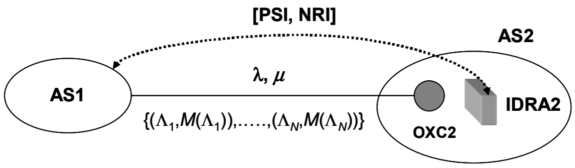

Fig.2. Estimation of the number of available wavelengths between two directly connected ASes, where the destination AS consists of only one OXC.

nario in Fig. 2 we assume the same arrival and departure rates $\forall\ \Lambda_i$. It is worth highlighting that in extended scenarios this might not always be the case. In Section II we explained the flexibility that an AS has while composing its NRI advertisements. Such flexibility can cause that a given destination is reachable for some wavelengths but unreachable for some others, so the traffic demands may differ for different wavelengths.

Let $x_i(t)$ denote the ENAW of type Λ_i at time t in the interdomain link in Fig. 2. Our goal is to estimate $x_i(t)$ according to the preceding traffic demands. Let $p_k(t)$ be the probability that the ENAW of type Λ_i at time t is k, that is, $p_k(t)=\text{Prob}\{x_i(t)=k\}$. The process $x_i(t)$ evolves in $t \in [(n-1)T, nT]$ according to the birth and death model in Fig. 3, where T denotes the observation time interval. This latter is the average time interval between two routing updates, taking into account the updates coming from NRI messages, the ones coming from PSI messages, and also the ones coming from the allocations performed by the local IDRA.

The birth and death process in Fig. 3 has $(M(\Lambda_i)+1)$ states, where the state "0" indicates that no wavelengths of type Λ_i are available in the interdomain link. Then, the state transitions can be described by the following set of differential equations for the probabilities $p_k(t)$:

$$\begin{aligned}\dot{p}_k(t) &= [\,M(\Lambda_i)-k+1\,]\mu p_{k-1}(t) - [\,k\lambda + (M(\Lambda_i)-k)\mu\,]\,p_k(t) \\ &\quad + (k+1)\lambda\, p_{k+1}(t)\end{aligned} \tag{5}$$

With boundary conditions:

$$\begin{aligned}\dot{p}_0(t) &= -M(\Lambda_i)\mu\, p_0(t) + \lambda\, p_1(t) \\ \dot{p}_{M(\Lambda_i)}(t) &= \mu\, p_{M(\Lambda_i)-1}(t) - M(\Lambda_i)\lambda\, p_{M(\Lambda_i)}(t)\end{aligned} \tag{6}$$

Then, we use the expected value of $x_i(t)$ as its estimator. Using (5) and (6), the expected value can be derived as follows:

$$\dot{E}[x_i(t)] = \frac{d}{dt}\sum_{k=0}^{M(\Lambda_i)} k p_k = \sum_{k=0}^{M(\Lambda_i)} k\, \dot{p}_k = \mu M(\Lambda_i) - (\lambda+\mu)E[x_i(t)] \tag{7}$$

Integrating (7) in the observation time interval yields (8):

$$E[x_i(nT)] = E[x_i((n-1)T)]e^{-(\lambda+\mu)T} + \left(\frac{\mu M(\Lambda_i)}{\lambda+\mu}\right)\left[1-e^{-(\lambda+\mu)T}\right] \tag{8}$$

Equation (8) allows to recurrently estimate $x_i(t)$. If the state is known at the beginning of the observation interval, (8) solves the estimation problem for the interdomain scenario in Fig. 2 until the next observation interval. As mentioned before extending this model for multiple domains and multiple traffic demands is not a trivial task. In the sequel we propose a straightforward way to roughly estimate $x_i(t)$ in such cases, and rely on the capabilities of a Kalman filter to refine this estimation.

Let l_{ID} denote a particular interdomain link of the AS for which we want to derive the estimation. We define $D_{\Lambda i}$ as the set of all possible destinations that are reachable through l_{ID} using wavelength Λ_i. We assume that the interdomain calls requesting a route through l_{ID} to a destination $d \in D_{\Lambda i}$ arrive as independent Poisson processes with exponentially distributed arrival rate λ_d. The duration of these calls are also assumed to be exponentially distributed with the same departure rate μ, and they are assumed to be independent of previous arrivals and holding times. Based on these assumptions, we propose to extend the estimation in (8) as follows:

$$E[x_i(nT)] \approx E[x_i((n-1)T)]e^{-\gamma_{i,n-1}T} + \left(\frac{\mu M(\Lambda_i)}{\gamma_{i,n-1}}\right)\left[1-e^{-\gamma_{i,n-1}T}\right] \tag{9}$$

$$\gamma_{i,n-1} = \left(\sum_{d\in D_{\Lambda_i}} \lambda_d \prod_{j=1}^{H^{adv}} \left[1-BR_j(d)\right] + \mu\right)_{t=(n-1)T} \tag{10}$$

The rationale in (9)-(10) is three-fold. First, the model captures the essential characteristics of state-dependent circuit switched networks. Second, the model is simple and easy to compute since it uses aggregated state information that is locally available. And third, the inherent coupling between the arrival rates and the occupancy is straightforwardly relaxed by means of the BR advertisements.

From (9) we define the following constants:

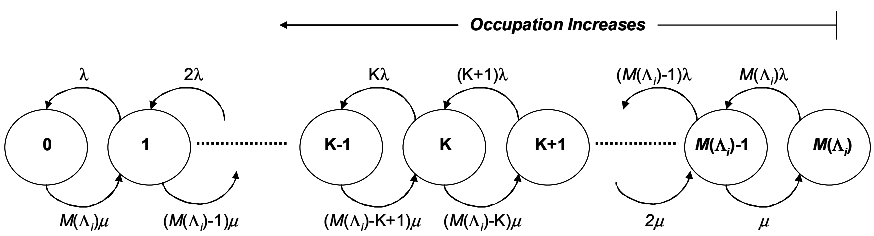

Fig.3. Birth and death process modeling the availability of wavelengths of type Λ_i for the two directly connected ASes.

$$A_{i,n-1} = e^{-\gamma_{i,n-1}T}$$
$$B_{i,n-1} = \left(\frac{\mu M(\Lambda_i)}{\gamma_{i,n-1}}\right)\left[1 - e^{-\gamma_{i,n-1}T}\right] \tag{11}$$

In order to simplify the notation we define $E[x_i(nT)]=x_{i,n}$. Then, from the approximation in (9):

$$\begin{cases} x_{i,n} = A_{i,n-1}x_{i,n-1} + B_{i,n-1} + w_{i,n-1} \\ w \sim (0,Q) \end{cases} \tag{12}$$

where w represents the process noise, which is assumed to be white, with zero mean and variance Q [4].

The linear stochastic difference equation in (12) is the main result of this section, and it is precisely the input for the discrete-time Kalman filter.

IV. Kalman Filter and the RWA Algorithm

Kalman filters are powerful tools since they offer a computationally efficient way to optimally estimate the state of a controlled process. The estimation is optimal in the sense that Kalman filters minimize the covariance of the estimation error [4]. The basic principle of Kalman filters is that they alternate between two steps, namely, a prediction step and a correction step. The idea is to predict the next state of a process based on the partial knowledge of the current state, and then adjust this prediction with the new information coming from the observations. The adjusted state is then considered as the new prediction and so on. In the sequel we introduce the prediction-correction steps for our particular problem. We start by defining the following set of variables in Table I.

TABLE I
Notation for the Kalman Filter

Symbol	Description
$x^-_{i,n}$	A *priori* estimate of the ENAW Λ_i on path P (predicted state)
$x^+_{i,n}$	A *posteriori* estimate of the ENAW Λ_i on path P (corrected state)
$e^+_{i,n} = x_{i,n} - x^+_{i,n}$	Estimation error
$\varepsilon^+_{i,n} = E[(e^+_{i,n})^2]$	Estimation error covariance
$K_{i,n}$	Kalman filters set their gain $K_{i,n}$ so as to minimize the estimation error covariance
$z_{i,n} = x_{i,n} + v_{i,n}$	$z_{i,n}$: is the observation. Denotes the ENAW Λ_i on a path P observed from the routing updates. $v_{i,n}$: is the observation noise, which is assumed to be white, with zero mean and variance R [4].

Using (12), the usual prediction-correction Kalman steps yield [4]:

Prediction Step:
$$x^-_{i,n} = A_{i,n-1}x^+_{i,n-1} + B_{i,n-1} \tag{13}$$
$$\varepsilon^-_{i,n} = A^2_{i,n-1}\varepsilon^+_{i,n-1} + Q$$

Correction Step:
$$K_{i,n} = \left[\frac{\varepsilon^-_{i,n}}{\varepsilon^-_{i,n} + R}\right] \tag{14}$$
$$x^+_{i,n} = x^-_{i,n} + K_{i,n}\left[z_{i,n} - x^-_{i,n}\right]$$
$$\varepsilon^+_{i,n} = \left[1 - K_{i,n}\right]\varepsilon^-_{i,n}$$

An important feature of Kalman filters is that its convergence is not biased by the initial state [4]. For the RWA algorithm we define a configurable threshold T_h that triggers the utilization of the filter in the RWA decision. When the effective number of *all* available wavelengths along the candidate paths is below or equal to T_h, the RWA is driven by the Kalman filter. If this is not the case the RWA is performed by choosing the ($P(s, d)$, Λ_i) pair with the lowest cost.

V. Simulation Results

In order to evaluate the Kalman-based estimation proposed in this paper, we have contrasted its performance against two different interdomain RWA algorithms, namely, OBGP+ and Cost. The former is the result of a set of enhancements that we introduced to Optical BGP (OBGP) [9], and which we call OBGP+. Similarly as BGP, OBGP is a shortest AS-path routing algorithm that exchanges reachability information, but it does not handle path state information. Our OBGP+ handles the wavelength availability information introduced in Section II. In this sense, OBGP+ learns and advertises the ENAW along the candidate paths. The RWA algorithm of OBGP+ is essentially a shortest AS-path least loaded routing algorithm. On the other hand, the Cost RWA algorithm is based on the additive cost introduced in Section II. This RWA handles both, wavelength availability information and load information. Thus, it is expected that it performs better than OBGP+ since it handles more path state information. Our results in Fig. 4 (a) confirm this fact. The Kalman-based estimation is basically a module aiding the Cost RWA. In the extreme case that the threshold T_h = 0, the Kalman RWA algorithm is off all the time, and the RWA algorithm is identical to the Cost algorithm.

As a first step toward the validation of our proposals, we have conducted a series of simulations in a multi-domain topology consisting of five domains with multiple connections between them. The topology is the same that we used in [10]. Each domain has several nodes inside it, and there are at least two disjoint paths between any given pair of nodes in each domain. We have used 5 fibers per-link, and 24 wavelengths per-fiber. The threshold for all the trials was set to T_h = 2. During the trials we have used different observation time intervals T, and this is shown in Fig. 4 for the cases T=1, T=3, and T=8 units of time. Our simulations were conducted using cross Poisson traffic between different domains, and as shown in Fig. 4, the trials were performed for different traffic loads.

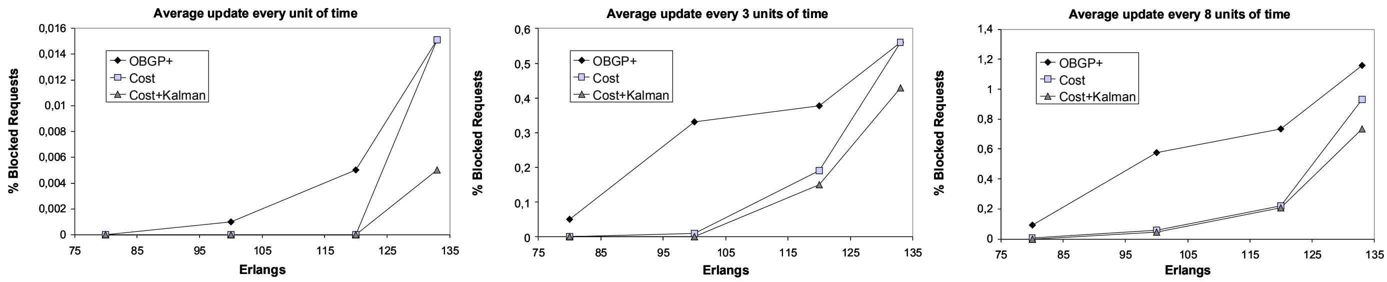

(a)

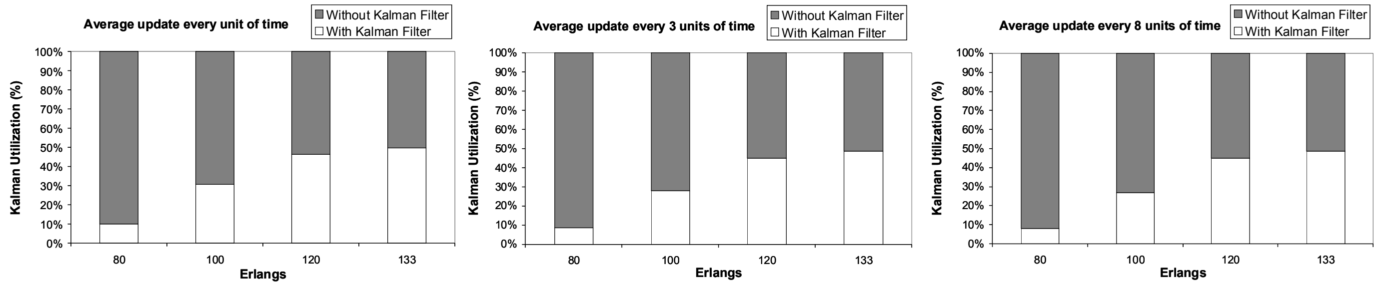

(b)

Fig. 4 (a) Percentage of blocked requests with OBGP+, Cost, and Cost+Kalman. (b) Percentage of RWA decisions that were taken using Kalman.

As a performance indicator of the different RWA algorithms, we have used the percentage of blocked requests, which is shown in Fig. 4 (a), for different traffic loads and for different observation time intervals. Clearly, the Kalman-based estimation outperforms OBGP+, and it supplies significant improvements when contrasted with the Cost RWA algorithm. As expected, these improvements are especially noticeable as traffic load in the network increases. In this case, the wavelength availability decreases and the Kalman estimation aids the Cost RWA algorithm. Conversely, when the traffic load is low the Kalman filter is barely used and hence the performance of the Cost and Kalman RWA algorithms are practically the same. Figure 4 (b) shows the percentage of RWA decisions that were taken using Kalman for the different traffic loads. The next table summarizes the relative percentage of improvement in the blocked requests for the highest simulated traffic load in the network, i.e. 133 Erlangs.

TABLE II
RELATIVE PERCENTAGE OF IMPROVEMENT IN THE BLOCKED REQUESTS FOR THE HIGHEST SIMULATED TRAFFIC LOAD (133 ERLANGS)

Observation time interval	% of improvement vs. OBGP+	% of improvement vs. Cost
T=1	67 %	67 %
T=3	23%	23%
T=8	37%	21%

VI. DISCUSSION AND FUTURE WORK

In this paper we have demonstrated the usefulness of developing simple models to roughly estimate the occupancy in multi-domain wavelength-routed optical networks, and then refine this estimation by predictive techniques. Our main results and conclusions apply to a rather small multi-domain optical network, so further research is needed to analyze the feasibility of this kind of approach in a large-scale environment. We consider that estimation techniques like the one proposed in this paper offer a promising line of work to address the trade-off between obtaining a low blocking ratio, and keeping the path state information as limited as possible.

REFERENCES

[1] G. Liu, C. Ji, V. Chan, "On the Scalability of Network Management Information for Inter-Domain Light-Path Assessment," IEEE/ACM ToN, Vol. 13, No. 1, February 2005.
[2] M. Yannuzzi, X. Masip-Bruin, and O. Bonaventure, "Open Issues in Interdomain Routing: a survey," IEEE Network, Vol. 19, No. 6, November/December. 2005.
[3] M. Yannuzzi, S. Sánchez-López, X. Masip-Bruin, J. Solé-Pareta and J. Domingo-Pascual, "A Combined Intra-Domain and Inter-Domain QoS Routing Model for Optical Networks," IFIP/IEEE ONDM 2005, Milan, Italy, February 2005.
[4] H. W. Sorenson (ed.), "Kalman Filtering: Theory and Application," IEEE Press, 1985.
[5] A. Kolarov, A. Atai, and J. Hui, "Application of Kalman Filter in High-Speed Networks," in Proceedings of IEEE Globecom'94, San Francisco, USA, November 1994.
[6] T. Anjali, C. Scoglio, J. de Oliveira, "New MPLS Network Management Techniques Based on Adaptive Learning," IEEE Transactions on Neural Networks, Vol. 16(5), pp. 1242-1255, September 2005.
[7] A. Sridhar an, K. N. Sivarajan, "Blocking in All-Optical Networks," IEEE/ACM ToN, Vol. 12, No. 2, April 2004.
[8] T. K. Nayak, K. N. Sivarajan, "A New Approach to Dimensioning Optical Networks," IEEE JSAC, Vol. 20, No. 1, January 2002.
[9] L. Wang, H. Zhang, X. Zheng, et al, "A Novel OBGP–based mechanism for Lightpath Establishment in WDM Mesh Networks", in Proceedings of ECOC 2003, Rimini, Italy, September 2003.
[10] E. Marín-Tordera, X. Masip-Bruin, S. Sánchez-López, J. Solé-Pareta, and J. Domingo-Pascual, "A Hierarchical Routing Approach for Optical Transport Networks," Computer Networks, Vol. 50, No. 2, February 2006.