Reactive Robust Routing: Anomaly Localization and Routing Reconfiguration for Dynamic Networks

Pedro Casas · Lionel Fillatre · Sandrine Vaton · Igor Nikiforov

Abstract This paper presents a novel approach to deal with dynamic and highly uncertain traffic in dynamic network scenarios. The Reactive Robust Routing (RRR) approach is introduced, a combination of proactive and reactive techniques to improve network efficiency and robustness, simplifying network operation. RRR optimizes routing for normal-operation traffic, using a time-varying extension of the already established Robust Routing technique that outperforms the stable approach. To deal with anomalous and unexpected traffic variations, RRR uses a fast anomaly detection and localization algorithm that rapidly detects and localizes abrupt changes in traffic flows, permitting an accurate routing adaptation. This algorithm presents well-established optimality properties in terms of detection/localization rates and localization delay, which allows for generalization of results, independently of particular evaluations. The algorithm is based on a novel parsimonious model for traffic demands which allows for detection of anomalies using easily available aggregated-traffic measurements, reducing the overheads of data collection.

Keywords: Traffic uncertainty Robust routing and load balancing Traffic modeling Network monitoring Optimal anomaly detection and localization

P. Casas (✉)
Universidad de la República, J. Herrera y Reissig 565, CP 11300 Montevideo, Uruguay
e-mail: pcasas@fing.edu.uy

L. Fillatre · I. Nikiforov
ICD/LM2S, Université de Technologie de Troyes, 12 rue Marie Curie, BP 2060, 10010 Troyes Cedex, France
e-mail: lionel.fillatre@utt.fr

I. Nikiforov
e-mail: igor.nikiforov@utt.fr

S. Vaton
Télécom Bretagne, Technopôle Brest-Iroise, CS 83818, 29238 Brest Cedex 3, France e-mail: sandrine.vaton@telecom-bretagne.eu

1 Introduction

The performance of the Internet itself depends, in large measure, on the operation of the underlying routing protocols. Today's routing protocols in IP networks compute routing configurations based on network topology and some rough knowledge of traffic demands (e.g., worst-case traffic, average traffic, long-term forecasts), without regard to current traffic load on routers or even possible traffic misbehaviors that often arise. Over-provisioning has been so far the main reason of success for such a "naive" approach. However, routing optimization is becoming increasingly difficult due to the dynamic nature and the increasing uncertainty of current traffic demands. The overall IP traffic is expected to grow from 6.6 exabytes per month in 2007 to nearly 29 exabytes per month by 2011, more than quadrupling in less than a half decade [1]. Simultaneously, the evolution of access technologies and the development of optical access networks (e.g. Fiber To The Home technology) is dramatically increasing the bandwidth for end-users, imposing serious and unforeseen problems at the core network, so far assumed infinitely provisioned. In this near future scenario, traditional over-provisioning may no longer be an economically viable solution to handle dynamic traffic demands.

In this paper we address the problem of routing optimization, automatic routing reconfiguration, and load balancing in a single Autonomous System (AS). The traffic demand in an AS network is typically described by a traffic matrix (TM), which captures the amount of traffic transmitted between every pair of ingress and egress nodes of the network, so called the Origin-Destination (OD) traffic flows. OD flows present two different behaviors: on the one hand, a stable and predictable behavior due to normal usage patterns (e.g. daily traffic demand fluctuation); on the other hand, a highly dynamic and unpredictable behavior due to unexpected events, such as network equipment failures, flash crowd occurrences, security threats (e.g. denial of service attacks), external routing changes (e.g. inter-AS routing through BGP), and new spontaneous overlay services (e.g. P2P applications for real-time content delivery). We shall use the term *volume anomaly* [2] to describe these unexpected events, characterized by large and sudden link load changes.

Recent works [3–5] have proposed a novel solution to the routing optimization under traffic uncertainty problem: the Stable Robust Routing (SRR) approach. In a robust perspective of routing optimization, demand uncertainty is taken into account directly within the optimization problem, computing a single routing configuration for all demands within some *uncertainty set*. While this routing configuration is not optimal for any single TM within the set, it minimizes the worst case performance over the whole set. In this sense, SRR provides worst-case performance bounds for all possible traffic variations within the uncertainty set. The obtained robust routing configuration is usually applied during long periods of time (e.g. daily routing), avoiding the potential instabilities induced by routing changes. The SRR approach can be therefore seen as a *proactive* technique to deal with dynamic traffic. It can

handle variable traffic demands up to a certain limit, given by the size of the uncertainty set. However, using a SRR algorithm to address both normal-operation traffic as well as volume anomalies is an inefficient strategy: it is clear that a single routing configuration cannot be adequate for both situations. On the contrary, a *reactive* approach could be used as a complementary strategy to enhance the SRR performance, responding to abrupt and large traffic modifications with an effective routing adaptation. Volume anomalies may have an important impact on network performance, causing sudden situations of strong congestion. The early detection and localization of these anomalies allows to adjust routing as soon as possible, limiting their impact. Unfortunately, there are several shortcomings in current algorithms for anomaly detection in the TM that limit their usefulness in the practice: (1) most detection methods rely on highly tuned data-driven traffic models that are not stable in time, hence they are not appropriate for the task; (2) current detection methods present a lack of theoretical support on their optimality properties, making it almost impossible to compare their performances; (3) most approaches only treat the detection of the anomaly, but the anomaly localization and the application of countermeasures are still open problems; just detecting the anomaly does not solve the real problem.

1.1 Related Work

There is a large literature on routing optimization for uncertain traffic demands. Traditional algorithms rely on a small group of TMs to compute quasi-optimal routing configurations. An extreme case is presented in [6], where routing is optimized for a single estimated TM and it is then applied for daily routing. Traffic uncertainty is characterized by multiple TMs in [7], and different mechanisms to find optimal routes for these TMs are presented. Given the dynamic nature of current traffic demands, the traditional perspective may no longer be suitable for current scenario [8]. A different perspective is provided by Dynamic Load Balancing (DLB) algorithms: TeXCP [9] and MATE [10] both balance load in real-time, responding to instantaneous traffic variations. Their main goal is to avoid network congestion by adaptively balancing the load among fixed paths, based on measurements. DLB presents a desirable property, that of keeping routing adapted to traffic variations. However, as we shall see, these adaptive algorithms may present poor performance under significant and abrupt traffic changes. A third category of algorithms consists in Stable Robust Routing (SRR) techniques [3–5]. In [3], authors capture traffic variations by introducing a polyhedral set of demands, applying linear programming techniques to compute an optimal single routing configuration for all the traffic demands within this set. Oblivious Routing [4] also defines linear algorithms to optimize worst-case performance for different sizes of traffic uncertainty sets. The drawback of SRR is its inherent dependence on the definition of the uncertainty set: larger sets allow to handle more variable traffic demands, but at the cost of routing inefficiency; tighter sets produce more efficient routings, but are subject to poor performance guarantees. In [5], authors introduce an approach to deal with this trade off in the size of the uncertainty set, combining traditional algorithms with oblivious routing. Routing is optimized for expected

traffic, and bounds are provided for worst-case performance in the event of large traffic variations. Nevertheless, this approach proposes a single routing configuration as previous works do, losing the adaptability, and hence the performance efficiency of dynamic algorithms. We shall evidence that it is possible not only to ensure performance guarantees for unexpected events, but to obtain near-optimal routing configurations to deal with this traffic.

As regards network anomaly detection, the problem has been extensively studied. We just overview the most representative works for network-wide anomaly detection in the TM, using aggregated link traffic measurements as input, SNMP measurements from now on [2–12]. In [2], authors use Principal Components Analysis (PCA) and the sub-space technique to detect and localize network-wide anomalies in the TM, separating the SNMP measurements into anomalous and anomaly-free traffic representations. The PCA approach is a data-driven method which highly depends on the particular set of SNMP measurements under analysis, and as shown in [13], it must be highly tuned to provide accurate results, becoming impractical. [11] analyzes SNMP measurements using ARIMA modeling, Fourier transforms, wavelets, and PCA to model traffic evolution, detecting and locating anomalies as large deviations. Unfortunately, the technique presents a lack of theoretical results on its optimality properties (e.g. maximization of the probability to detect an anomaly with a bounded false alarm rate), a necessary condition to provide robust and easy-to-generalize results. [12] uses a Kalman-filtering approach to track the evolution of OD flows from SNMP measurements, detecting anomalies as large prediction errors. The method requires a long training phase where direct anomaly-free OD flow measurements are used to calibrate the underlying model. As we have recently shown [14], the assumed model has a particular structure that may require several periodical re-calibrations to provide reliable results, which makes it too costly to implement in the practice. Besides, the paper does not tackle the anomaly localization problem.

1.2 Contributions of the Paper

This work proposes both proactive and reactive complementary approaches to manage dynamic traffic demands, separately treating both sources of traffic variation. For *expected traffic variations*, we present a time varying extension of SRR that outperforms current single static-routing approach: Multi-Hour Robust Routing (MHRR). SRR may be costly: however, it is easy to control its cost by shrinking the uncertainty set. In MHRR, the uncertainty set is optimally divided into several sub-sets, considering the time direction for the partitioning. MHRR preserves the virtues of SRR, but changes the routing configuration during time, applying a SRR configuration for each period of the day. For the case of *unexpected traffic variations*, we present an optimal algorithm for volume anomaly detection and localization, permitting to identify strong traffic variations, locate their origins, and perform accurate routing changes. To overcome the stability problems of previous anomaly detection approaches, a novel linear, parsimonious, non data-driven traffic model is proposed. The model is used to filter the anomaly-free traffic from the SNMP measurements, providing residuals sensitive to anomalies. The

main contribution of the detection/localization algorithm relies on its well-established optimality properties, a fundamental feature generally absent in previous works.

Both proactive and reactive methods are combined into a novel approach to handle dynamic traffic demands: Reactive Robust Routing (RRR). Figure 1 presents a high level description of this approach. RRR uses MHRR to handle normal variations in traffic demands, and the detection/localization algorithm to deal with unexpected volume anomalies. RRR exploits the localization ability to deploy an adapted robust routing configuration after the anomalous traffic detection, reducing its impact on network performance during its prevalence. In addition, it also provides a simple yet effective method to automatically detect the end of the anomaly, regaining the MHRR configuration. Contrary to previous works in the field, this proposal optimizes routing in a robust and adaptive fashion for every possible traffic demand, and not only for the common-case or expected traffic. A key feature of RRR relies on the fact that the whole routing modification algorithm is completely automatic, an interesting property that simplifies network operation. The proposed algorithms are validated using real traffic from three different backbone networks: the Internet2 Abilene backbone network, the European GEANT network, and a private international Tier-2 network. This work represents a continuation of our previous works on robust routing [8, 15] and anomaly detection [16].

The remainder of this paper is organized as follows. In Sect. 2, the basic concepts of the robust routing approach are recalled. Section 3 presents the theoretical background and empirical evaluation of MHRR. The linear parsimonious TM model is introduced and validated in Sect. 4. Section 5 describes the anomaly detection/localization algorithm and presents an empirical validation using real traffic. Section 6 presents the Reactive Robust Routing approach, showing the automatic interaction between the proactive and the reactive components through complete real and simulated anomaly scenarios. In this section we also propose a load-balancing

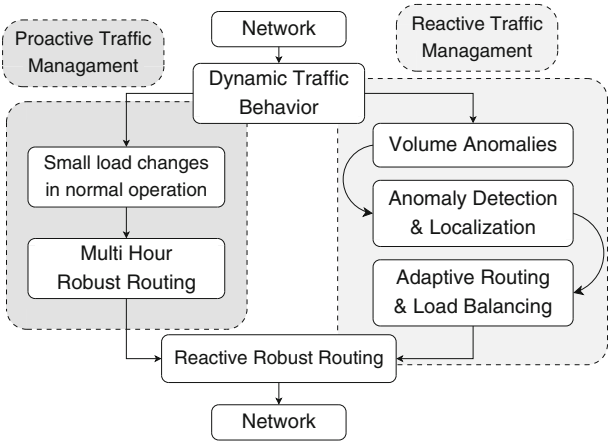

Fig. 1 High level description of the Reactive Robust Routing approach

extension for RRR nicknamed Reactive Robust Load Balancing (RRLB), in which end-to-end paths remain fixed and only routing fractions are modified. We further compare this extension against traditional DLB methods, showing the impressive stability of RRLB when faced with volume anomalies. Finally, Sect. 7 concludes this work.

2 Robust Routing

Let us consider a network topology defined by a set of n nodes and a set of r links $L = \{1, ..., r\}$ with capacities in $C = \{c_1, c_2, ..., c_r\}$. The TM $\mathbf{d} = \{d(i, j)\}$ denotes the traffic flow between every node i and node $j (i \neq j)$ of the network. This matrix is arranged as a column vector, $\mathbf{d} = \{d(k)_{,\ k=1..m}\}$, where $d(k)$ represents the Origin-Destination traffic flow transmitted by each OD pair of nodes k, and $m = n\ .\ (n - 1)$ is the number of OD pairs. Let $N = \{\mathrm{OD}_1, ..., \mathrm{OD}_m\}$ be the set of m OD pairs. Let $y(l)$ be the total aggregated traffic at link l in a certain period of time. This data is available from routers management information and it is periodically collected via the well-known SNMP protocol. Traffic demands and links traffic are related through the routing matrix R, a $r \times m$ matrix $R = \{r_{l,k}\}$ where $0 \leqslant r_{l,k} \leqslant 1$ represents the fraction of OD flow k routed through link l: $\mathbf{y} = R\ .\ \mathbf{d}$, with $\mathbf{y} = \{y(l)_{,l=1..r}\}$. Routing optimization depends on the underlying data transport mechanism; in this case, the focus is on path-based routing such as MPLS. This optimization consists of minimizing certain performance index associated with traffic demands and network topology. Throughout this work, the Maximum Link Utilization (MLU) is considered as the routing performance criterion. For a given routing matrix R and a traffic demand $\mathbf{d}$, the MLU $u_{\max}$ is defined as the maximum of the ratio between link traffic and link capacity (1):

$$u_{\max}(C, \mathbf{d}, R) = \max_{l \in \{1...r\}} \sum_{k=1}^{m} \frac{r_{l,k} \cdot d(k)}{c_l} = \max_{l \in \{1...r\}} \frac{y(l)}{c_l} \tag{1}$$

While this criterion does not represent a direct measure of Quality of Service (QoS), overloaded links tend to cause QoS degradation (e.g. larger delays and packet losses, throughput reduction, etc.), so MLU represents a reasonable measure of network performance. There are many other performance indexes that could be used instead of MLU, like end-to-end path delay or mean link utilization; setting the focus too strictly on the MLU can often lead to longer average traffic paths and thus adversely affect the mean. However, the MLU is by far the most commonly applied criterion as it represents an easy to understand performance measure, so it will be the adopted one. Let $P(k)$ be the set of possible paths for OD flow k. Let $r_{p,k}$ be the proportion of traffic demand $d(k)$ that flows through path $p \in P(k)$, $0 \leqslant r_{p,k} \leqslant 1$. Link routing fractions $r_{l,k}$ can be directly computed from path routing fractions $r_{p,k}$:

$$r_{l,k} = \sum_{p \in P(k),\, p \ni l} r_{p,k} \qquad \forall k \in N, \forall l \in L \tag{2}$$

$$
\begin{aligned}
&\text{minimize} \quad u_{\max} \\
&\text{subject to:} \\
&\sum_{p \in P(k)} r_{p,k} = 1 \qquad \forall k \in N \\
&\sum_{p \in P(k),\, p \ni l} r_{p,k} = r_{l,k} \qquad \forall k \in N,\, \forall l \in L \\
&\sum_{k \in N} r_{l,k}.d(k) \leqslant u_{\max}.c_l \qquad \forall l \in L,\ \forall \mathbf{d} \in D \\
&r_{p,k}, r_{l,k} \geqslant 0 \qquad \forall l \in L, \forall p \in P(k),\, \forall k \in N \\
&u_{\max} \leqslant 1
\end{aligned}
\tag{3}
$$

Routing optimization for traffic demand $\mathbf{d}$ can be easily performed when this demand is perfectly known. However, in the practice, $\mathbf{d}$ is unknown and all we may know for sure is that it belongs to a certain bounded set D in which it can vary. This set represents the *uncertainty* in the value of $\mathbf{d}$, and so it is usually known as the *uncertainty set*. The set D can be defined in different ways, depending on the available data: SNMP measurements and historical routing, a set of previously measured TMs, time series of TMs $\mathbf{d}_t$, etc. Authors in [3] define this set as a *polytope*, based on the intersection of several half-spaces that results from linear constraints imposed to traffic demands. As a practical example, we shall define a simple uncertainty set D, based on the routing matrix R and the busy-hour links traffic $\mathbf{y}_{\text{busy}}$:

$$
D = \left\{ \mathbf{d} \in \mathbb{R}^m, R.\mathbf{d} \leqslant \mathbf{y}_{\text{busy}},\, \mathbf{d} \geqslant 0 \right\} \tag{4}
$$

Figure 2a depicts the obtained polytope D, based on the convex intersection of r half-spaces $\mathbf{r}_i \cdot \mathbf{d} \leqslant y_{\text{busy}}(i)$, $\forall i \in \{1 \ldots r\}$, where $\mathbf{r}_i$ stands for the i-th row of R.

The Robust Routing Optimization Problem (RROP) defined in (3) consists of minimizing $u_{\max}$ for all the traffic demands within the uncertainty set D. This linear system can be efficiently solved by linear programming techniques, applying a combined columns and constraints generation method [3]. RROP represents a worst-case optimization for all the traffic demands inside D, hence it provides performance guarantees $\forall \mathbf{d} \in D : u_{\max}(C, \mathbf{d}, R_{\text{robust}}) \leqslant u^*_{\max}$, $\forall \mathbf{d} \in D$, where $u^*_{\max}$ and $R_{\text{robust}} = \left\{ r^*_{l,k} \right\}$ are the solution to (3). Note that using (4) as uncertainty set has

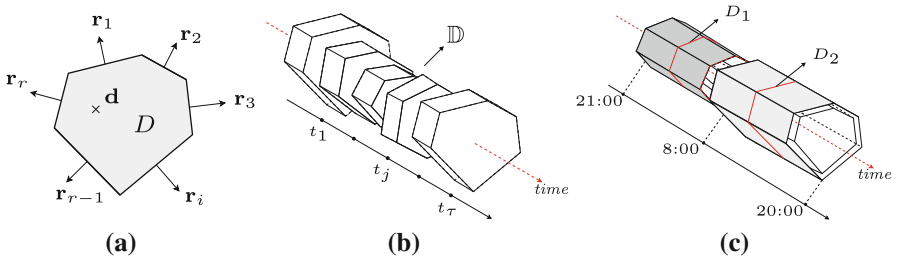

Fig. 2 **a** The set D as a polytope, **b** daily polytope $\mathbb{D}$, **c** time partitioning of $\mathbb{D}$

a major advantage: routing optimization can be performed from easily available SNMP measurements, without even measuring any TM at all. In a traditional robust routing application, the obtained routing configuration R_{robust} is applied during long periods of time, usually in a daily routing basis. In this sense, the robust routing approach is usually referred to as Stable Robust Routing (SRR).

3 Multi-Hour Robust Routing

The advantages of SRR with respect to traditional routing approaches are presented in [8]; to sum up, SRR offers stability guarantees against traffic uncertainty and normal traffic variations at a reasonable cost. However, considering a single routing configuration for long periods of time is not a cost-effective solution and results in sub-optimal performance. We therefore propose an approach to shrink and adapt the uncertainty set that outperforms SRR. Based on rough knowledge of traffic variations, basically considering expected traffic behavior, the uncertainty set is optimally divided in the direction of the time, producing several sub-sets; a multi-hour routing configuration is then built, considering a single SRR configuration for each of these sub-sets.

Daily traffic variations can be seen as a time variation of the uncertainty set. At each time slot t_j, the routing matrix R and the SNMP measurements $\mathbf{y}_{t\_j}$ define an uncertainty set $D(t_j) = \left\{\mathbf{d} \in \mathbb{R}^m, R.\mathbf{d} \leqslant \mathbf{y}_{t_j}, \mathbf{d} \geqslant 0\right\}$. The slotted time comes from the fact that SNMP measurements $\mathbf{y}_{t_j}$ are collected at discrete time intervals t_j, typically every 5 or 10 mins. The union of several uncertainty sets along contiguous time slots $t_1, .., t_j, .., t_\tau$ defines a *temporal uncertainty set* $\mathbb{D} = \left\{\mathbf{x} = \left\{\mathbf{d}_{t_j}, t_j\right\} \in \mathbb{R}^{m+1}, \mathbf{d}_{t_j} \in D(t_j), t_j \in [t_1..t_\tau]\right\}$. Figure 2b explains this idea. Assuming that this set is a union of polytopes, [17] provides a theoretical study of the optimal partitioning of $\mathbb{D}$, using a partitioning hyper plane. In particular, it proves that this is a NP-hard problem, except for the case where a partitioning direction is previously fixed. In such a case, the author presents a simple algorithm to approximately solve (5) in polynomial-time, using a generalization of a simple dichotomy methodology.

A partitioning hyper plane is defined by its direction vector α and a value β:α. $\mathbf{x} = \beta$. In MHRR there is a particular direction for partitioning: the *time direction*. In this case, $\alpha = [0, .., 0, 1]$ and $\beta = t$. Given β_1 and β_{h+1}, we define $h - 1$ hyperplanes at times $\boldsymbol{\beta} = \{\beta_2, .., \beta_h\}$. Let D_i be the convex hull of the union $\left\{\cup_{\beta_i \leqslant t_j \leqslant \beta_{i+1}} D(t_j)\right\}, \forall i = 1, ., h$, see Fig. 2c. MHRR consists in computing the optimal times $\boldsymbol{\beta}^*$ when routing should be modified, in order to minimize the worst case of MLU in $\mathbb{D}$; $\boldsymbol{\beta}^*$ is the solution to (5), where $u^*_{\max}(D_i)$ is the solution to (3) for polytope D_i.

$$\boldsymbol{\beta}^*(\mathbb{D}) = \arg\min_{\boldsymbol{\beta}} \left\{ \max_{i=1..h} u^*_{\max}(D_i) \right\} \tag{5}$$

Finally, a single SRR configuration is computed for each time interval $\left[\beta^*_i, \beta^*_{i+1}\right], \forall i = 1, ., h$. The interesting issue in our proposal is that we provide an

objective means of computing an optimal multi-hour routing design, maintaining the robustness of the SRR approach. The optimality property of our approach lies on the computation of the ideal times $\boldsymbol{\beta}^*$ to switch routing. Traditional methods used in the design of multi-hour routing configurations are rather simplistic, relying on a couple of TMs to optimize different routing configurations [18].

MHRR presents a trade-off between performance and routing stability. The more intervals, the more adapted the routing becomes. However, the number of intervals should be bounded as many routing changes may lead to instabilities and performance degradation. In a general case, 2 sub-sets are enough to handle the usual daily variation [8]. Additionally, and as we shall explain in Sect. 6.3, it is possible to modify just the routing fractions $r_{p,k}$ and keep the same paths $P(k)$ used in each SRR, avoiding routing reconfigurations. In other words, it is possible to do Multi-Hour Load Balancing instead of MHRR.

3.1 MHRR Evaluation

The SRR and the MHRR approaches are compared in Abilene, an Internet2 backbone network. Abilene consists of 12 router-level nodes and 30 links. The used network topology and traffic demands are available at [19]. Traffic data consists of 6 months of TMs collected via Netflow from the Abilene Observatory [20] in 2004. As the measured traffic demands do not significantly load the network, we have rescaled them by multiplying all their entries by a constant factor. Since the Abilene network spans different time zones, the time variation of the polytope is not a simple homothety: in fact, there is no strong traffic synchronization between links. In this sense, a routing scheme modification during the day improves routing performance. Let R_o be the historical routing matrix used in Abilene, which is not necessarily optimal. This routing matrix is also available at [19]. A single time partitioning is considered (i.e. 2 routing intervals), $\beta_1 = 21{:}00$, $\beta_2 = \beta^*$, and $\beta_3 = 20{:}00$, where β^* is the solution to (5). The smallest polytope that includes all possible realizations over that period is computed for each time interval:

$$D_1 = \{\mathbf{d} \in \mathbb{R}^m, R_o.\mathbf{d} \leqslant \mathbf{y}_1,\ \mathbf{d} \geqslant 0\} \tag{6}$$

$$D_2 = \{\mathbf{d} \in \mathbb{R}^m, R_o.\mathbf{d} \leqslant \mathbf{y}_2,\ \mathbf{d} \geqslant 0\} \tag{7}$$

where $\mathbf{y}_1 = \mathbf{y}^{\max}_{21:00-\beta^*}$ and $\mathbf{y}_2 = \mathbf{y}^{\max}_{\beta^*-20:00}$ include the maximum values for each link in the corresponding period. In this way, D_1 includes all the traffic demands between 21:00 and β^*, and D_2 between β^* and 20:00, see Fig. 2c. For each polytope, a SRR configuration is computed, R^1_{robust} and R^2_{robust}. In order to compare the stable and the multi-hour approaches, both routing configurations are applied during the whole evaluation period. The routing performance obtained with R_o is also included, which corresponds to the dotted line with label *Historical Routing* in Fig. 3. Figure 3a compares the MLU obtained by these two SRR configurations. Polytope D_1 is better suited for smaller loads, so R^1_{robust} performs better during the first half of the day, when network load is lower. However, when traffic increases, demands that do not belong to D_1 produce higher link utilizations than those obtained with R^2_{robust}. MHRR consists of computing the optimal time when routing must be changed,

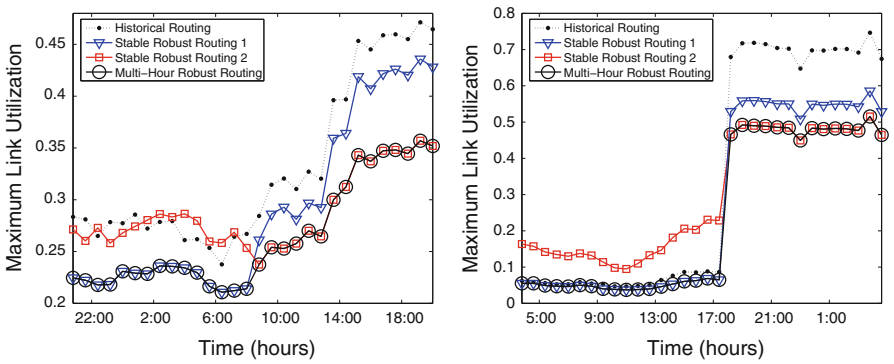

(a) Expected daily behavior **(b)** Anomalous unexpected event

Fig. 3 Routing performance, Stable versus Multi-Hour Robust Routing. **a** Traffic corresponds to TMs from the 1st week of the dataset, starting on the 01/03/04. **b** TMs from the 23th week, starting on the 28/08/04

using the corresponding routing configuration depending on the time of the day: R^1_{robust} before β^* and R^2_{robust} after. In this evaluation, the computed value of $\beta^* \approx$8:00. MHRR presents a performance improvement of 16% w.r.t. SRR before β^*, reaching a near 20% of over-efficiency after β^*. Traffic demands that drastically change due to a volume anomaly are considered as a second case study. Figure 3b presents an huge abrupt change in MLU at time 18:00. In this case, we shall assume that the change is known in advance; note that in the general case it is not possible to predict such abrupt variations. The optimal time for switching routing is $\beta^* \approx$ 18:00. MHRR definitely outperforms SRR in this traffic scenario, providing a MLU between 10 and 60% smaller during the evaluation period.

4 Dealing with Unexpected Events

The proposed MHRR approach offers a robust and efficient routing configuration, given a rough knowledge of the temporal uncertainty set. However, in the presence of volume anomalies it is no longer possible to apply MHRR; the optimal division proposed in (5) cannot be done if $\mathbb{D}$ presents strong and unknown variations. A reactive approach is proposed for those cases, based on the detection and localization of these volume anomalies. To avoid costly and difficult to perform direct OD flow measurements, the detection and localization algorithm uses SNMP measurements $\mathbf{y}_t$ as input data. Additional flow-based technology is necessary to network-wide collect and process direct OD flow measurements [21], so different traffic models have been developed using SNMP measurements and routing information to *reconstruct* OD flows. This reconstruction represents an ill-posed problem, because the number of unknown OD flows is much larger than the number of links [21]; briefly, it is not possible to compute $\mathbf{d}_t = \{d_t(1), ..., d_t(m)\}$ from $\mathbf{y}_t = R \,.\, \mathbf{d}_t$, because $r \ll m$. To overcome this problem, a novel parsimonious, linear model for normal-operation traffic is proposed next. In this work, this traffic model

is used to remove the anomaly-free traffic from the anomaly detection problem. However, the model can be used to solve other problems, like the well-known TM estimation problem among others.

4.1 Stochastic Traffic Model for Anomaly Detection

We assume that the stochastic process of the TM $\mathbf{d}_t$ obeys the following linear expression:

$$\mathbf{d}_t = \boldsymbol{\lambda}_t + \boldsymbol{\xi}_t \tag{8}$$

where $\boldsymbol{\lambda}_t \in \mathbb{R}^m$ is the mean traffic demand and $\boldsymbol{\xi}_t$ is a white Gaussian noise with covariance matrix $\Sigma = \text{diag}\left(\sigma_1^2, \ldots, \sigma_m^2\right)$. The process $\boldsymbol{\lambda}_t$ represents the temporal evolution of the mean TM, which can be correctly modeled in the absence of anomalies. The white Gaussian noise $\boldsymbol{\xi}_t$ models the natural variability of the TM together with the modeling errors. In order to describe $\boldsymbol{\lambda}_t$ with a small number of coefficients, a key feature of the TM is employed: its spatial stationarity. Many classical TM models make use of this assumption, like the very well-known gravity model [22, 23]. The other key observation for this model is the *mice and elephants phenomenon*: a small percentage of OD flows contribute to a large proportion of the total traffic in the TM [22]. The existence of such dominant flows together with the spatial stationarity of flows makes it reasonable to assume that, in the absence of an anomaly, the largest OD flows in a network remain the largest, and the smallest flows remain the smallest during long periods of time; this assumption is confirmed in the empirical validation of the model, at least for several days, see Sect. 4.3. Therefore, it seems logical to accept that the order of increasing OD flows w.r.t. their traffic volume remains stable in time. The sorted OD flows can be interpreted as a discrete non-decreasing signal with certain smoothness. The curve obtained by interpolating this signal is assumed to be a continuous curve, hence it can be parameterized by a polynomial splines approximation.

Figure 4 shows the anomaly-free OD flows for three different operational networks, sorted in the increasing order of their volume of traffic, for different time instants t. The dashed lines depict the value of each sorted OD flow $d_t(k)$, $k = 1..m$,

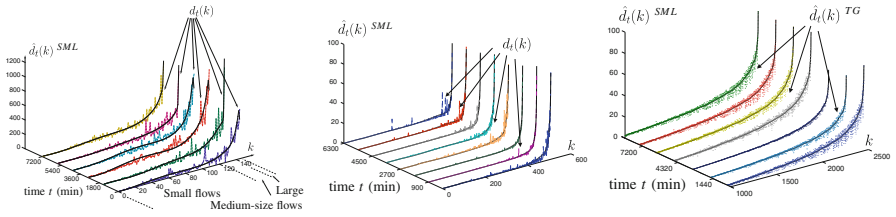

(a) Abilene Network **(b)** GEANT Network **(c)** Tier-2 ISP Network

Fig. 4 Approximation of real OD flows (*dashed lines*) by the spline-based model (*full lines*) in 3 operational networks. $\hat{d}_t(k)^{TG}$ is the estimated OD flow k using the tomogravity estimation method, introduced in [23]. TMs come from the week starting on the **a** 01/03/04, **b** 01/06/05, and **c** 18/04/05, respectively

the full lines represent the polynomial approximation of the sorted flows. In order to appreciate the time stability of this approximation, the curves are plotted for several consecutive days. Given the shape of the curve formed by the sorted OD flows, a cubic splines approximation is applied; basic definitions and results on polynomial splines can be found in [24]. A discrete spline basis is designed, discretizing the continuous splines according to m points uniformly chosen in the interval [1;m] and rearranging them according to the OD flows sorting order. The obtained linear parsimonious model for the anomaly-free traffic demand can be expressed as:

$$\mathbf{d}_t = S\boldsymbol{\mu}_t + \boldsymbol{\xi}_t \tag{9}$$

where $S = \{\mathbf{s}(i),\ i = 1..q\}$ is a $m \times q$ known matrix with a small number of columns w.r.t. the number of OD flows, i.e., $q \ll m$. The vectors $\mathbf{s}(i)$, which correspond to the rearranged discrete splines, form a set of known basis vectors describing the spatial distribution of traffic. The vector $\boldsymbol{\mu}_t = \{\mu_t(1)\ldots\mu_t(q)\}^T$ is the unknown time varying parameter that describes the OD flow intensity distribution w.r.t. the set of vectors $\mathbf{s}(i)$. The model for the anomaly-free links traffic is given by:

$$\mathbf{y}_t = G\boldsymbol{\mu}_t + \boldsymbol{\zeta}_t, \tag{10}$$

where $G = RS$ and $\boldsymbol{\zeta}_t \sim \mathcal{N}(0, \Phi)$, with $\Phi = R\Sigma R^T$. The computation of the rank of G is not simple since it depends on R. In the practice, since the number of columns of G is very small, the product RS and its rank can be computed very fast. Therefore, it will be assumed that G is full column rank. To simplify notation and computations, the *whitened* measurements vector is introduced:

$$\mathbf{z}_t = \Phi^{-\frac{1}{2}}\mathbf{y}_t = H\boldsymbol{\mu}_t + \epsilon_t, \tag{11}$$

where $H = \Phi^{-\frac{1}{2}}G$ and $\epsilon_t \sim \mathcal{N}(0, I_r)$ (I_r is the $r \times r$ identity matrix). The purpose of this transformation is simply to whiten the Gaussian noise. Finally, the covariance matrix Σ is unknown. The solution consists in computing a simple empirical estimate $\widehat{\Sigma}$ from a few anomaly-free measurements.

4.2 Validation of the Model: The Dataset

The validation of the proposed traffic model is conducted using real data from two operational networks: the Abilene network previously presented, and the European GEANT research network. GEANT traffic data consists of 15' sampled TMs, built from IGP and BGP routing information and Netflow data, available at the TOTEM website [25]. In the following evaluations, we assume that traffic demands $\mathbf{d}_t$ are unknown and just consider the link load values $\mathbf{y}_t$ as the known data.

In order to verify the stability properties of the model, two sets of measurements are used: the first one, the *learning* anomaly-free dataset, is composed of one hour of anomaly-free SNMP measurements (3 h in GEANT, due to the different sampling rates) and it is used to construct the spline basis S; the second one, the *testing* dataset, is composed of 672 SNMP measurements and it is used to validate the model. Let T_{learning} and T_{testing} be the set of time indexes associated with SNMP

measurements from the learning and the testings datasets, respectively. The learning dataset is measured immediately before the testing dataset.

4.3 Numerical Validation of the Model

In order to validate the splines model, we shall consider an estimate of the real traffic demands and test its accuracy. We shall construct a Maximum-Likelihood (ML) estimate for $\mathbf{d}_t$. The statistical properties of the ML estimate are very well-known, and thus it represents an excellent estimation approach. Since the traffic model is a Gaussian model, the Spline-based Maximum Likelihood (SML) estimate $\hat{\mathbf{d}}_t^{SML}$ corresponds to a simple least-squares estimate:

$$\hat{\mathbf{d}}_t^{SML} = \left(S\left(H^T H\right)^{-1} H^T \Phi^{-\frac{1}{2}}\right)\mathbf{y}_t \tag{12}$$

The Spline-based model is computed using the learning dataset, following these steps: (1) a tomogravity estimate $\hat{d}_t^{TG}(k)$ [23] is computed for all OD flows k and all $t \in T_{\text{learning}}$, (2) the mean OD flow values $\bar{d}^{TG}(k)$ are computed: $\bar{d}^{TG}(k) = \frac{1}{\#(T_{\text{learning}})} \sum_t \hat{d}_t^{TG}(k)$, and (3) they are sorted in ascending order to obtain a rough estimate of the OD flows traffic volume. The Spline-based model is designed with cubic splines and 2 knots, representing small, medium-size, and large OD flows. The estimates $\hat{d}_t^{TG}(k)$ and mean values $\bar{d}^{TG}(k)$ are also used to compute an estimate $\hat{\sigma}_k^2$ of σ_k^2, which leads to an estimate $\hat{\Phi}$ of Φ.

As a global indication of the accuracy of the SML estimate, and to test its performance against some other well-known traffic estimation methods, we shall use the relative Root Mean Squared Error (RRMSE) for each time t of the testing dataset:

$$\text{RRMSE}(t) = \frac{\sqrt{\sum_{k=1}^{m} \left(d_t(k) - \hat{d}_t^{\text{label}}(k)\right)^2}}{\sqrt{\sum_{k=1}^{m} d_t(k)^2}}, \forall t \in T_{\text{testing}} \tag{13}$$

where $d_t(k)$ is the true traffic volume of the anomaly-free OD flow k at time t and $\hat{d}_t^{\text{label}}(k)$ denotes the corresponding estimate for the method entitled 'label'. The RRMSE provides at each time t a summary of the relative estimation error for the TM. Figure 5a depicts the CDF of the RRMSE errors for the 672 TMs in Abilene and GEANT, showing that in more than 70% of the time indexes, estimation errors are below 10%. A deeper study of the RRMSE shows that in most cases, large RRMSE values correspond to large relative errors in the smallest-volume OD flows, which are well known to be hard to estimate [23, 26]. Note however that small OD flows have little impact on traffic engineering tasks, hence are generally less important to estimate. The mean values of the RRMSE for the evaluation period are 8.14% in Abilene and 7.04% in GEANT. Methods proposed in the literature as "accurate" estimates present relative errors that vary between 5 and 15% [26], so obtained results are satisfactory.

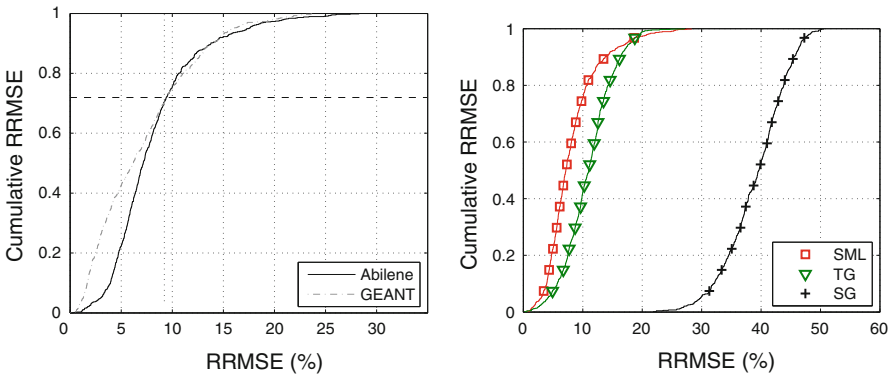

(a) CDF - RRMSE for the SML estimates. **(b)** CDF - RRMSE in Abilene.

Fig. 5 Comparison between the SG, TG and SML estimates for 672 anomaly-free TMs from the week starting on the 01/03/04 in Abilene, and on the 01/06/05 in GEANT

Figure 5b compares in Abilene the performance of the SML estimate with two very well-known methods for TM estimation: the simple gravity and the tomogravity estimates [23], with labels 'SG' and 'TG', respectively. The RRMSE corresponding to the tomogravity estimate TG is quite close to the error produced with our model. However, the SML estimate presents a major advantage w.r.t. the TG estimate: as it was previously said, the ML estimate presents well established statistical properties, which is not the case for the TG estimate. The SML estimate is *asymptotically optimal*, i.e., it is asymptotically unbiased and efficient. Moreover, the Spline-based model can be used in order to design anomaly detection algorithms with optimality properties, which is not the case for the tomogravity estimate.

As a final validation, the Gaussian assumption of the model is studied. The *residuals* of measurements are analyzed, i.e. the obtained traffic after filtering the "regular" part $H\boldsymbol{\mu}_t$. The residuals are obtained by projection of the whitened measurements vector $\mathbf{z}_t = \Phi^{-\frac{1}{2}}\mathbf{y}_t$ onto the left null space of H, obtaining a residuals vector $\mathbf{u}_t = W\mathbf{z}_t \sim \mathcal{N}(0, I_{r-q})$, see the Appendix for computation details about the *rejection matrix* W. A simple Kolmogorov-Smirnov test at the level 5% accepts the Gaussian hypothesis for 662 of the 672 anomaly-free measurements, representing an acceptation ratio of 98.5%. This confirms the Gaussian assumption of the model.

5 Anomaly Detection and Localization

Once the traffic model has been introduced and validated, the reactive component of our proposal is presented. The goal of the proposed method is to detect and localize an additive change $\theta\ \mathbf{r}_j$ in a series of SNMP measurements, where $\mathbf{r}_j$ is the j-st column of the routing matrix R, and θ is the intensity of the anomaly; this corresponds to a change θ in OD flow j. In this work we consider the same simplifying hypothesis as in [2], considering only "localized" anomalies, namely anomalies in a single OD flow at a time. As it is shown in [2, 27], this hypothesis is

extensively verified in both the Abilene network and the Sprint-Europe network, the European backbone of a US tier-1 ISP. For example, [27] shows that more than 85% of the volume anomalies detected in Abilene correspond to single-OD flow anomalies. The simplifying hypothesis is therefore adequate in real-traffic scenarios.

We propose an optimal sequential algorithm to detect abrupt changes in traffic demands. Sequential approaches are used to minimize the number of observations needed to detect an anomaly after its occurrence. The sequential algorithm not only detects an abrupt change in the TM, but also localizes the change, i.e. it identifies the OD flow in which the change has occurred. This detection/localization algorithm is optimal in the sense that it minimizes the maximum mean detection/localization delay for a given maximum probability of false localization η and minimum mean time before a false alarm υ, a usual measure of the false alarm rate. The localization of the anomalous traffic is possible since an anomaly in a given OD flow typically spans multiple links, and has a particular *signature* as a result of the routing process.

The detection/localization of a volume anomaly that occurs at an unknown time t_0 can be treated as a hypothesis testing problem, where the null hypothesis $\mathcal{H}_0 = \{\text{OD flows are anomaly-free}\}(t_0 = +\infty)$ is tested against m alternatives $\mathcal{H}_{t_0}^j = \{$the j-th OD flow presents an anomalous additional amount of traffic θ from time $t_0\}$, $j = 1..m$. The algorithm has to compute the alarm time T_r at which a v_r-type change $\in \{1, 2, ..., m\}$ is detected and localized, based on SNMP measurements $\mathbf{y}_1$, $\mathbf{y}_2$, ... The hypothesis testing can be written as:

$$\mathcal{H}_0 : \mathbf{z}_t \sim \mathcal{N}(H\boldsymbol{\mu}_t, I_r),\ t = 1, 2, \ldots, \tag{14}$$

$$\mathcal{H}_{t_0}^j : \begin{cases} \mathbf{z}_t \sim \mathcal{N}(H\boldsymbol{\mu}_t, I_r),\ t = 1, \ldots, t_0 - 1, \\ \mathbf{z}_t \sim \mathcal{N}(H\boldsymbol{\mu}_t + \theta\Phi^{-\frac{1}{2}}\mathbf{r}_j, I_r), \\ \quad \theta_{\text{lower}} \leqslant |\theta| \leqslant \theta_{\text{upper}},\ t = t_0, t_0 + 1, \ldots \end{cases} \tag{15}$$

where $\mathbf{z}_t$ is the whitened measurements vector and $0 < \theta_{\text{lower}} < \theta_{\text{upper}} < +\infty$ are some known bounds on the change intensity of the j-th OD flow that are introduced for technical reasons, but can be arbitrarily chosen. Hypothesis (15) can be rewritten by filtering the anomaly-free traffic (see the Appendix):

$$\mathcal{H}_{t_0}^j : \begin{cases} \mathbf{u}_t \sim \mathcal{N}(0, I_{r-q}),\ t = 1, \ldots, t_0 - 1, \\ \mathbf{u}_t \sim \mathcal{N}(\theta\mathbf{v}_j, I_{r-q}), \\ \quad \theta_{\text{lower}} \leqslant |\theta| \leqslant \theta_{\text{upper}},\ t = t_0, t_0 + 1, \ldots \end{cases} \tag{16}$$

where $\mathbf{u}_t = W\mathbf{z}_t$ are the residuals obtained from $\mathbf{z}_t$ after filtering the anomaly-free traffic, $\mathbf{v}_j = W\Phi^{-\frac{1}{2}}\mathbf{r}_j$ is a known non-null vector and the matrix W is the linear rejector that eliminates the anomaly-free traffic (see the Appendix). The vector $\mathbf{v}_j$ corresponds to the signature in the residuals of a change in OD flow j, specified by the column $\mathbf{r}_j$ of R.

The recursive algorithm proposed in [28] perfectly fits this detection and localization problem, with one ideal feature, that of minimizing the number of samples needed to decide between the hypotheses with bounded false localization and false alarm rates. The algorithm produces two outputs: a stopping time T_r that

represents the time when the anomaly is detected and a decision v_r about which OD flow carries the anomaly:

$$T_r = \min_{1 \leqslant k \leqslant m} \{T_r(k)\}, \quad v_r = \arg \min_{1 \leqslant k \leqslant m} \{T_r(k)\} \tag{17}$$

$$T_r(k) = \inf \left\{ t \geqslant 1 : \min_{0 \leqslant j \neq k \leqslant m} [g_t(k,j) - h_{k,j}] \geqslant 0 \right\}, \; k = 1, \ldots, m \tag{18}$$

with $g_t(k, j) = g_t(k, 0) - g_t(j, 0)$. At each time t, the recursive functions $g_t(k,0)$ give an idea of the difference between the value of traffic residuals under the hypothesis of normal operation and the hypothesis of anomaly. Each recursive function $g_t(k,0)$ is defined as:

$$g_t(k,0) = (g_{t-1}(k,0) + u_t(k,0))^+ \tag{19}$$

$$u_t(k,0) = \log \frac{f_k(\mathbf{u}_t)}{f_0(\mathbf{u}_t)} \tag{20}$$

where $(x)^+ = \max(x, 0)$, $g_0(k, 0) = 0$ for every $1 \leqslant k \leqslant m$, and $g_t(0,0) = 0$ for all t. The function f_0 represents the probability density function of residuals in normal-operation conditions. The function f_k is the probability density function of residuals $\mathbf{u}_{t_0}, \mathbf{u}_{t_0+1}$, ... after a change of type k. The thresholds $h_{k,j}$ are chosen by the following formula:

$$h_{k,j} = \begin{cases} h_\mathrm{d} \text{ if } 1 \leqslant k \leqslant m \text{ and } j = 0 \\ h_\mathrm{i} \text{ if } 1 \leqslant k, j \leqslant m \text{ and } j \neq k \end{cases} \tag{21}$$

where h_d and h_i are the detection and localization thresholds. $T_r(k)$ is the first time when one of the recursive functions $g_t(k, j)$ exceeds the thresholds $h_{k,j}$. The detection time T_r corresponds to the earliest off all the times $T_r(k)$, $1 \leqslant k \leqslant m$. The detected anomaly is declared in OD flow j if the earliest of all these times was $T_r(j)$. The algorithm is asymptotically optimal, i.e. it reaches the lower bound of the maximum mean delay for detection/localization [28], given bounds in the mean time between false alarms υ and in the probability of false localization η:

$$\mathbb{E}_0(T_r) \geq \upsilon \tag{22}$$

$$\mathrm{Pr}_{t_0}^k(v_r = j | T_r \geq t_0) \leq \eta, \tag{23}$$

$\mathbb{E}_0(\cdot)$ denotes the expectation when all the measurements have the same probability density function f_0 (anomaly-free behavior) and $\mathrm{Pr}_{t_0}^k$ corresponds to the probability (with density function f_k) that the sequential test declares the final decision j whereas the true change type is $k \neq j$. The choice of the detection and localization thresholds h_d and h_i is discussed in [29].

5.1 Numerical Validation of the Anomaly Detection/Localization Algorithm

We demonstrate the ability of the sequential algorithm to detect and localize a volume anomaly from SNMP measurements in two different networks: a large Tier-

2 network (50 nodes, 168 links and 2,450 non-zero OD flows, sampled at a 10' rate) and Abilene. Figure 6a,b show the typical realization of the decision functions $g_t(i,0)$ and $s_t(i) = \min 0 \leqslant k \neq i \leqslant m \ [g_t(i,k) - h_{i,k}]$ in the Tier-2 network. Functions $s_t(i)$ are used to "monitor" the OD flows; when $s_t(i)$ exceeds the threshold 0, OD flow i is declared anomalous. The anomaly begins at time 3,660. Note that after this time, several decision functions $g_t(i,0)$ rapidly grow. Each function $g_t(i,0)$ is associated with OD flow i and when this function grows, it means that OD flow i is suspected of carrying an abnormal amount of traffic. Contrary to $g_t(i,0)$, only decision function $s_t(159)$ associated to OD flow 159 exceeds the localization threshold. Hence, functions $s_t(i)$ permit to localize the anomalous OD flow among all the OD flows associated to functions $g_t(i,0)$ that have rapidly grown. At time 3660, an alarm is raised and the algorithm selects the anomalous OD flow 159. The decision function $s_t(i)$ needs only 1 measurement to localize the anomalous OD flow. Similar results are obtained in Abilene in Fig. 6c,d, for a volume anomaly in OD flow 87.

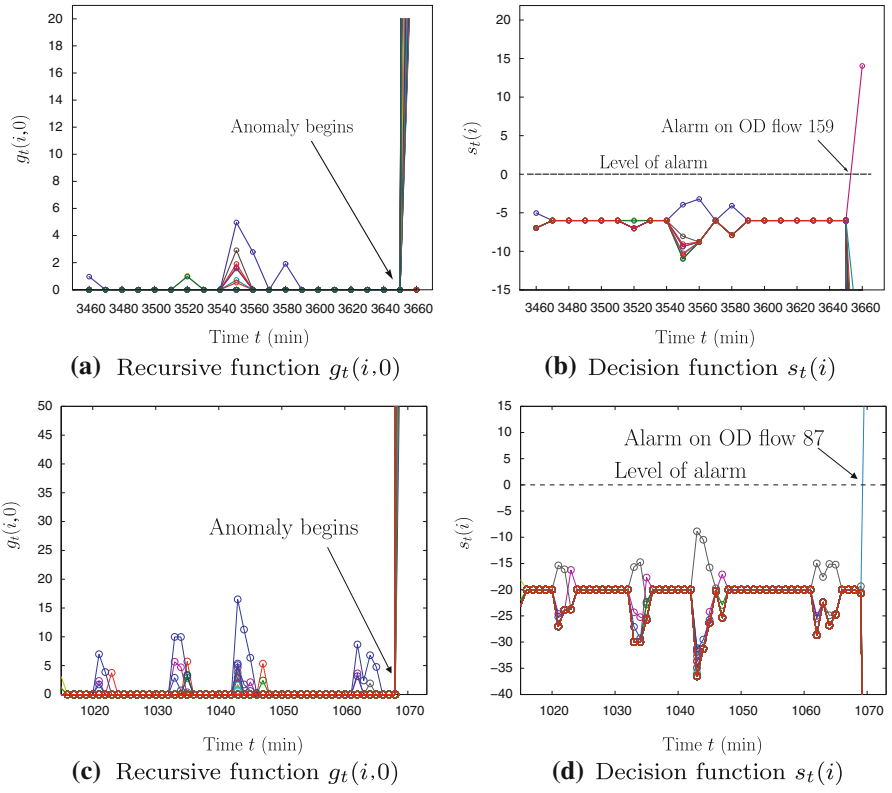

(a) Recursive function $g_t(i,0)$ **(b)** Decision function $s_t(i)$

(c) Recursive function $g_t(i,0)$ **(d)** Decision function $s_t(i)$

Fig. 6 Typical realizations of decision functions for **a**,**b** a Tier-2 network and **c**, **d** Abilene. TMs come from the week starting on the 01/05/04 in Abilene, and on the 04/04/05 in the Tier-2 network

6 Reactive Robust Routing

Both proactive and reactive methods, namely MHRR and the anomaly detection/localization algorithm, respectively, are combined into a single approach, introduced in this paper as Reactive Robust Routing (RRR). The reader is referred to the high level description of RRR in Fig. 1. This approach provides an automatic method for robust routing configuration/reconfiguration, based on the monitoring of the network state. RRR exploits the localization ability of the detection/localization algorithm to deploy a new and adapted robust routing configuration after the detection of an anomalous OD flow; at the same time, it detects the end of the anomaly (if there is any) and takes back the usual MHRR configuration.

6.1 Routing Reconfiguration

A simple method that exploits both the SRR approach and the localization ability of the detection/localization algorithm is proposed to compute a new routing configuration to handle a volume anomaly. The idea of this reconfiguration is to minimize the impacts of the detected anomaly on the network performance. As part of MHRR, we shall assume that before the detection of the anomalous traffic, a SRR configuration R_{robust} is used, computed on the basis of some historical routing R_{o} and some measured links traffic $\mathbf{y}_{\text{o}}$.

Let us consider an anomalous increase θ in the traffic volume of OD flow k. The normal-operation traffic demand $\mathbf{d}$ takes the value $\mathbf{d}_{\text{anomaly}} = \mathbf{d} + \boldsymbol{\theta}$, with $\boldsymbol{\theta} = \theta.\boldsymbol{\delta}_k$, where $\boldsymbol{\delta}_k = (\delta_{1,k}, \ldots, \delta_{k,k}, \ldots, \delta_{m,k})^T, \delta_{i,k} = 0$ if $i \neq k$ and $\delta_{i,i} = 1$. In RRR, the normal-operation uncertainty set D is expanded in the directions of the routed anomalous OD flow k, obtaining an expanded uncertainty set $D_{\text{anomaly}} = \left\{\mathbf{d}_{\text{anomaly}} \in \mathbb{R}^m, R_{\text{o}}.\mathbf{d}_{\text{anomaly}} \leqslant \mathbf{y}_{\text{o}} + \mathbf{r}_{\text{o}_k}\theta, \mathbf{d}_{\text{anomaly}} \geqslant 0\right\}$. The key issue of this expansion is that the new uncertainty set contains now the anomaly, hence a new robust routing configuration can be computed, with the great advantage of being highly adapted to properly handle this anomalous traffic. Figure 7 shows the expansion of the uncertainty set. In the figure, an anomaly of volume θ occurs in OD flow k, which should span over links i, j, and h according to R_{o}. Note that the expansion of the uncertainty set D is done in the directions provided by the k-th column $\mathbf{r}_{\text{o}_k}$ of R_{o}, and not in the directions provided by the real routing configuration R_{robust}. The reader should bear in mind that the anomalies that we deal with originate outside the network and propagate between origin-destination nodes (e.g. flash crowds, denial of service attacks) and therefore, they do not depend on the particular routing configuration used in the network of analysis; the expansion with respect to R_{o} can be thought as a proactive way of computing an uncertainty set that contains an anomaly in OD flow k.

The new SRR scheme $R_{\text{robust}}^{\text{anomaly}}$ is the solution to (3), using D_{anomaly} as the uncertainty set. To avoid the estimation of the unknown anomalous volume θ, D can be expanded to the limits of links capacity c_i (or some fraction λ chosen by the network operator), in the directions of OD flow k : $D_{\text{anomaly}} = \left\{\mathbf{d}_{\text{anomaly}} \in \mathbb{R}^m, R_{\text{o}}.\mathbf{d}_{\text{anomaly}} \leqslant \boldsymbol{\tau}\right\}$, where $\tau_i = y_{\text{o}}(i)$ if $r_{\text{o}_{i,k}} = 0$, and $\tau_i = \lambda c_i$ if $r_{\text{o}_{i,k}} > 0$. Note that the

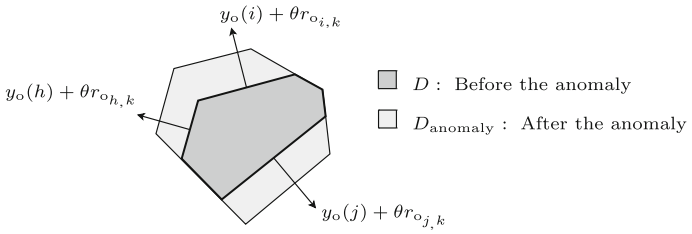

Fig. 7 Robust routing reconfiguration, based on the expansion of the uncertainty set

new routing configuration can be pre-computed off-line for each possible single-OD flow anomaly, obtaining a set of m SRR configurations $R_{\text{robust}}^{\text{anomaly}(k)}, k = 1, \ldots, m$.

Figure 8 presents the evaluation of RRR in the presence of a sudden and abrupt load change. This experiment considers the same situation depicted in Fig. 3b, and compares the routing performance of MHRR and RRR, respectively. Similar to Sect. 3, we assume that the daily uncertainty set is completely known for the case of MHRR (i.e. the abrupt change is known in advance). In RRR, the anomaly is automatically detected and the new routing configuration, computed a-priori on the basis of the expanded uncertainty set, is immediately deployed. The reader can appreciate that the routing performance of RRR is slightly worse than the one obtained by MHRR, less than 2%. Nevertheless, the RRR represents a real scenario, where the anomaly has to be detected in real-time to conduct an accurate routing reconfiguration.

6.2 Back to the MHRR Scenario

In order to regain the MHRR configuration after the end of the detected anomaly, RRR provides a method to detect the return to normal operation. This detection can

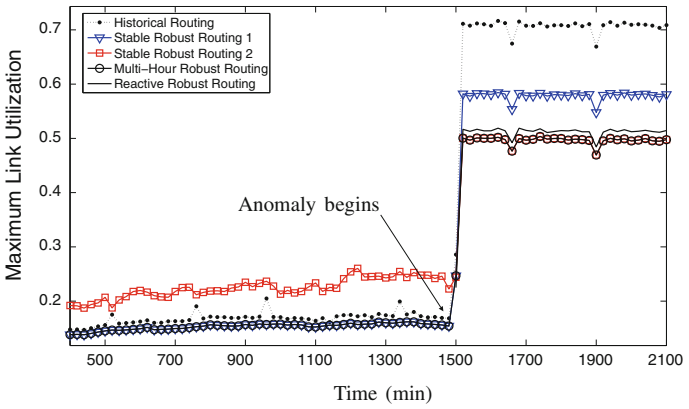

Fig. 8 Reactive Robust Routing—routing reconfiguration after the detection of a large and abrupt traffic change. TMs correspond to traffic demands from the 23th week of the Abilene dataset, starting on the 28/08/04

be easily achieved by using a similar hypothesis-testing approach to the one used to detect anomalies: suppose that a k_{type} anomaly has been detected and localized at time t_{anomaly}. For every time $t > t_{\text{anomaly}}$, we focus on the distribution of traffic residuals $f_k(\mathbf{u}_t)$ and look for a change that indicates the end of this anomaly. The reader should remember that in this work we have only considered anomalies in a single OD flow at a single time, and thus the only change that we can expect is back to normal operation.

Two simple hypotheses are considered for this problem, the null hypothesis $\mathcal{H}_0^k = \left\{\mathbf{u}_t \sim \mathcal{N}(\theta \mathbf{v}_k, I_{r-q})\right\}$, where the k-th OD flow presents an anomalous additional amount of traffic, against the alternative hypothesis $\mathcal{H}_{\text{alt}} = \left\{\mathbf{u}_t \sim \mathcal{N}(0, I_{r-q})\right\}$ where OD flow k is anomaly-free. A Neyman-Pearson test [30] is applied at each time t to decide between $\mathcal{H}_0^k$ and $\mathcal{H}_{\text{alt}}$. The Neyman-Pearson test represents the most powerful test for two simple hypotheses [30]. The statistics of this test is given by:

$$\Lambda(\mathbf{u}_t) = \log \frac{f_{\text{alt}}(\mathbf{u}_t)}{f_k(\mathbf{u}_t)} - h \gtrless 0 \tag{24}$$

where the decision threshold h is defined according to the desired false alarm probability. In this case, the function f_{alt} represents the probability density function of residuals under anomaly-free behavior, i.e hypothesis $\mathcal{H}_{\text{alt}}$, while f_k is the probability density function of residuals in the presence of an anomaly in OD flow k, i.e. hypothesis $\mathcal{H}_0^k$. If $\Lambda(\mathbf{u}_t) < 0$, the decision test chooses hypothesis $\mathcal{H}_0^k$. When $\Lambda(\mathbf{u}_t) > 0$, the test decides hypothesis $\mathcal{H}_{\text{alt}}$, pointing out the end of the anomaly.

To conclude this section, Fig. 9 presents an evaluation of the complete RRR approach under the presence of a simulated volume anomaly. An artificial, sudden and large volume change is introduced in OD flow 63 of the Abilene dataset. This artificial traffic is introduced on top of the usual daily traffic between times 1,125 and 1,350. The first step of RRR consists in computing the MHRR configurations, using an expected daily uncertainty set. The optimal division (5) results in $\beta^* = 1{,}230$. The evaluation begins at time 1,020, when MHRR decides to apply R^1_{robust} (SRR 1 in Fig. 9c). The detection/localization algorithm continuously monitors the network state, and at time $t_{\text{anomaly}} = 1{,}125$ detects and localizes an anomalous behavior in OD flow 63 (Fig. 9a). After the detection, and before the new sampling of link loads, the new routing configuration is deployed, which was previously computed using the notions introduced in Sect. 6.1. At time $t = t_{\text{anomaly}} + 1$ the new routing configuration is active, and the anomaly-end detection phase begins. It is important to note that the matrix $H = \Phi^{-\frac{1}{2}} RS$ as well as the anomaly-free traffic *rejector* W must be recomputed after the change of the routing matrix R; in fact, the same re-computation must be conducted every time the routing matrix changes, restarting the detection algorithm to avoid transient effects. The decision statistics $\Lambda(\mathbf{u}_t)$ remains negative for every time $t > t_{\text{anomaly}}$, until time $t' = 1350$, when the positive value of $\Lambda(\mathbf{u}_{t'})$ shows the end of the anomalous behavior in OD flow 63. At this time, RRR compares t' with β^* in order to decide which routing to apply, R^1_{robust} if $t' < \beta^*$ or R^2_{robust} if $t' > \beta^*$. Once the new routing configuration is established, the anomaly detection/localization algorithm starts again to recursively

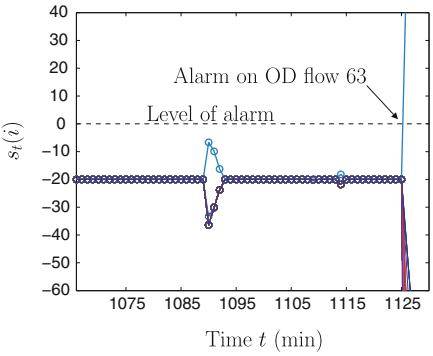

(a) Anomaly detected in OD flow 63.

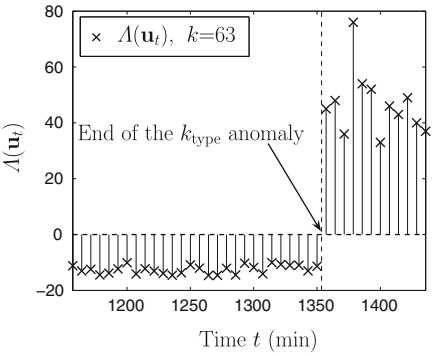

(b) Anomaly-end detected.

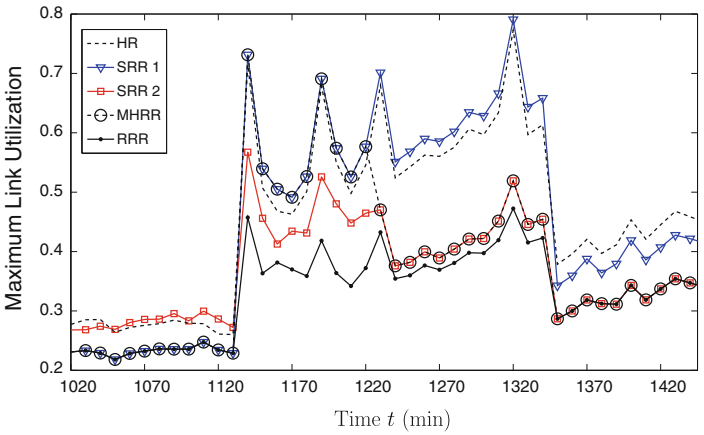

(c) Reactive Robust Routing.

Fig. 9 Reactive Robust Routing performance under a simulated volume anomaly. TMs correspond to traffic demands from the 12-th week of the Abilene dataset, starting on the 12/06/04

search for anomalies. The performance improvements of RRR are remarkable, up to 40% w.r.t. MHRR and near 50% w.r.t. the traditional SRR approach.

6.3 Reactive Robust Load Balancing versus Dynamic Load Balancing

Both MHRR and RRR consist in routing reconfiguration. However, even partially modifying the routing configuration of a large-scale network is in the practice a challenging task. It it easy to see that all our routing optimization algorithms can be very easily extended to the case of load balancing, where paths remain always fixed and the only modification is related to the fractions of traffic sent through each of these paths. Such a solution can be easily implemented in the practice, using path-based protocols such as MPLS. We believe that results presented so far will not significantly vary in case of load balancing, mainly because path diversity is rich enough and path modifications were usually rare in the presented evaluations. To

confirm this, we present and evaluate a load balancing extension of RRR, which we shall nickname Reactive Robust Load Balancing (RRLB). The RRLB approach uses a fixed set of paths $P(k)$, $\forall\ k \in N$, and only modifies the routing fractions $r_{p,k}$. Note that such a routing configuration can be directly computed from (3) using $P(k)$ fixed, applying a simpler constrains generation method instead of a combined columns and constraints generation method. The set of paths $P(k)$ used by RRLB is the one obtained from (3) in the computation of the SRR configuration for normal-operation traffic.

Let us evaluate the performance obtained by RRLB when faced with volume anomalies; additionally, we shall compare its execution with the one obtained by a traditional Dynamic Load Balancing algorithm, known as TeXCP [9]. In TeXCP, a convex link congestion function is defined in terms of link capacity and link load. The objective is to minimize the MLU, which is achieved by minimizing the biggest utilization that each OD flow obtains in its paths. This minimization is done recursively by means of a greedy algorithm, in which each single OD flow increases the amount of traffic sent along the path with the smallest utilization, independently of the other $m - 1$ OD flows. Such an adaptive algorithm is periodically executed, based on feedback from the network utilization. This constitutes the most challenging aspect of DLB, because convergence speed might be over-killing, specially under large and abrupt changes in traffic demands.

Figure 10 depicts the MLU obtained by RRLB and TeXCP in the traffic scenario presented in Fig. 8. To be as fair as possible, both algorithms use the same set of paths. In this evaluation, TeXCP adapts routing fractions $r_{p,k}$ every minute, meaning that for each new TM, five updates will be performed (recall that TMs are collected every 5' in Abilene). Results are shown then for every minute. The curve with label 'Actual Minimum' corresponds to the optimal value of MLU, computed for each single TM. Let us begin by TeXCP. A first important observation is that TeXCP has an important overshoot that causes high congestion, with an absolute difference w.r.t. the optimum of about 40%. The convergence after the anomaly is very slow, taking more than 6 h. However, it should be noted that when it eventually

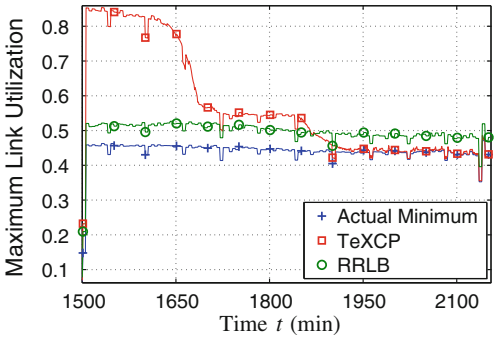

Fig. 10 Reactive Robust Load Balancing (RRLB) versus Dynamic Load Balancing (DLB). Transient behavior of DLB might be over-killing. RRLB avoids transient congestion thanks to the proactive computation of robust load balancing configurations. TMs correspond to traffic demands from the 23th week of the Abilene dataset, starting on the 28/08/04

converges, it obtains an optimal MLU. On the contrary, RRLB does not present a transient congestion behavior, basically because the routing fractions are not continuously modified but precomputed a-priori and immediately applied after the detection of the anomaly. Note that RRLB does not reach optimality, but the performance degradation is only 5%. This simple evaluation shows the great advantage of our approach, which provides a load-balancing technique without highly disruptive transient behaviors. Network operators are reluctant to use dynamic mechanisms mainly because they are afraid of this kind of transient behavior, and as we have seen, these concerns are not without reason. However, RRLB imposes as a plausible solution to the problem.

6.4 Numerical Complexity of RRLB and Implementation Issues

Data collection overhead, numerical complexity, and memory usage are central issues to determine the extra overhead introduced by RRLB. We shall focus on the reactive part of RRLB, namely the anomaly detection/localization algorithm, because the computation of the different SRR configurations used by RRLB is performed off-line, hence it does not limit its applicability.

Data collection is not really an issue in RRLB, for two main reasons: firstly, our algorithms use highly aggregated measurements, namely links traffic measurements, sampled at relatively low rates, e.g., one set of measurements every 5 min in this work, hence the cost of data collection is almost negligible; additionally, SNMP measurements are available in virtually every IP device, and are usually collected in every operational network, thus measurements are readily there to be used, without necessity of additional monitoring technology.

As regards anomaly detection, the method stores two matrices in memory, the "whitening" matrix $\Phi^{-\frac{1}{2}} \in \mathbb{R}^{r\times r}$, and the rejector $W \in \mathbb{R}^{(r-q)\times r}$. Given the recursive execution of the method, m additional variables are kept in memory, corresponding to the m recursive functions $g_t(i,0)$. For anomaly localization purposes, the m anomaly signatures $\mathbf{v}_k \in \mathbb{R}^{(r-q)\times 1}$ are also stored. This represents a total of $\mathcal{O}(r^2)$ variables in memory, where r is the number of links in the network. The computation of $\Phi^{-\frac{1}{2}}$ and W involves matrix multiplications and inversions, and thus the associated cost is $\mathcal{O}(r^3)$. There is an additional cost in the learning phase of the splines-based model, related to the tomogravity estimate used to construct the splines basis S. The cost of the tomogravity method is similar to that of the least-squares method, which implies $\mathcal{O}(m^3)$ operations to estimate an $m \times 1$ TM. All these matrices are computed off-line during the learning phase and do not affect the scalability and on-line applicability of the method. In the on-line detection/localization phase, residuals $\mathbf{u}_t = W\mathbf{z}_t$ are computed and used to update the m recursive functions $g_t(i,0)$. Finally, the m functions $s_t(i)$ used for anomaly localization are computed. These steps involve approximately $\mathcal{O}(r^2)$ operations for anomaly detection and $\mathcal{O}(m^2)$ additional operations for anomaly localization.

To sum up, RRLB keeps in memory a total of $\mathcal{O}(r^2)$ variables and implies $\mathcal{O}(r^2 + m^2)$ operations for on-line anomaly execution. This analysis shows that the overheads of RRLB are quite low, and demonstrates that the method is easily

scalable with the size of the network. All previous SRR approaches do not incur the overheads of RRLB, but they certainly provide higher MLU values or even become unfeasible in the event of volume anomalies (i.e. $u_{\max} > 1$), simply because there is no single routing configuration that can efficiently manage all possible traffic variations; it is clear from our study that some form of dynamism is necessary when faced with highly dynamic traffic.

7 Conclusions

In this paper, we addressed the routing optimization under traffic uncertainty problem. We introduced the Reactive Robust Routing approach, a proactive and reactive solution that not only deals with current dynamic traffic demands in a robust and efficient way, but also detects and localizes volume anomalies from aggregated links traffic measurements, improving network operation. We extended the robust routing technique to manage expected traffic variations by introducing the notion of time-varying uncertainty set, setting up a multi-hour robust routing scheme. This approach achieves better resource utilization than previous stable robust proposals in different scenarios. We introduced an original linear model to parameterize usual traffic behavior from widely available SNMP measurements. Compared to other traffic models, ours is not data-driven, a necessary property to achieve reliable results. In order to manage large and unexpected traffic variations, we presented a statistical algorithm to detect and localize volume anomalies in network traffic, using the proposed model to filter the anomaly-free traffic. This algorithm presents well-established optimality properties, which are extremely important to provide solid results. We proposed a novel approach that exploits both the Robust Routing paradigm and the localization ability of our detection/localization algorithm to optimally handle anomalous traffic variations. Contrary to traditional Dynamic Load Balancing algorithms, the load balancing extension of RRR does not present transient congestion behaviors, a necessary condition to apply dynamic routing in the practice.

Appendix: Elimination of the Anomaly-free Traffic

The anomaly-free traffic $H\boldsymbol{\mu}_t$ is removed by projecting the whitened measurements vector $\mathbf{z}_t = \Phi^{-\frac{1}{2}}\mathbf{y}_t$ onto the left null space of H (i.e. $WH = 0$). Let us define the matrix $W^T = (\mathbf{w}_1, .., \mathbf{w}_{r\text{-}q})$ of size $r \times (r - q)$, composed of eigenvectors $\mathbf{w}_1, .., \mathbf{w}_{r\text{-}q}$ of the projection matrix $P_H^{\perp} = I_r - H(H^TH)^{-1}H^T$, corresponding to eigenvalue 1. The matrix W satisfies the following conditions: $WH = 0, W^TW = P_H^{\perp}$, and $WW^T = I_{r\text{-}q}$. The matrix W can be considered as a linear rejector that eliminates the anomaly-free traffic. In the presence of an anomaly in OD flow j, traffic residuals $\mathbf{u}_t = W\mathbf{z}_t$ can be modeled as $\mathbf{u}_t = \theta W\Phi^{-\frac{1}{2}}\mathbf{r}_j + W\epsilon_t$. By defining the vectors $\mathbf{v}_j = W\Phi^{-\frac{1}{2}}\mathbf{r}_j$, we get the following distribution for traffic residuals in the presence of an anomaly in OD flow j : $\mathbf{u}_t \sim \mathcal{N}(\theta\mathbf{v}_j, I_{r-q})$.

References

1. Cisco Systems: Global IP Traffic Forecast and Methodology, 2006–2011. http://www.hbtf.org/files/cisco_IPforecast.pdf. Accessed 05 Aug 2010
2. Lakhina, A., Crovella, M., Diot, C.: Diagnosing Network-Wide Traffic Anomalies. Proceedings of the SIGCOMM'04 (2004). doi:10.1145/1015467.1015492
3. Ben-Ameur, W., Kerivin, H.: Routing of uncertain traffic demands. Optim. Eng. **6**, 283–313 (2005)
4. Applegate, D., Cohen, E.: Making Intra-Domain Routing Robust to Changing and Uncertain Traffic Demands: Understanding Fundamental Tradeoffs. Proceedings of the SIGCOMM'03 (2003). doi:10.1145/863955.863991
5. Wang, H., Xie, H., Qiu, L., Yang, Y., Zhang, Y., Greenberg, A.: COPE: Traffic Engineering in Dynamic Networks. Proceedings of the SIGCOMM'06 (2006). doi:10.1145/1151659.1159926
6. Roughan, M., Thorup, M., Zhang, Y.: Traffic Engineering with Estimated Traffic Matrices. Proceedings of the IMC'03 (2003). doi:10.1145/948205.948237
7. Zhang, C., Liu, Y., Gong, W., Kurose, J., Moll, R., Towsley, D.: On Optimal Routing with Multiple Traffic Matrices. Proceedings of the INFOCOM'05 (2005). doi:10.1109/INFCOM.2005.1497927
8. Casas, P., Vaton, S.: An Adaptive Multi Temporal Approach for Robust Routing. Euro-FGI Workshop on IP QoS and Traffic Control, IST Press, Portugal (2007)
9. Kandula, S., Katabi, D., Davie, B., Charny, A.: Walking the Tightrope: Responsive yet Stable Traffic Engineering. Proceedings of the SIGCOMM'05 (2005). doi:10.1145/1090191.1080122
10. Elwalid, A., Jin, C., Low, S., Widjaja, I.: MATE: MPLS Adaptive Traffic Engineering. Proceedings of the INFOCOM'01 (2001). doi:10.1109/INFCOM.2001.916625
11. Zhang, Y., Ge, Z., Greenberg, A., Roughan, M.: Network Anomography. Proceedings of the IMC'05, 317–330 (2005)
12. Soule, A., Salamatian, K., Taft, N,: Combining Filtering and Statistical Methods for Anomaly Detection. Proceedings of the IMC'05, 331–344 (2005)
13. Ringberg, H., Soule, A., Rexford, J., Diot, C.: Sensitivity of PCA for Traffic Anomaly Detection. Proceedings of the SIGMETRICS'07 (2007). doi:10.1145/1254882.1254895
14. Casas, P., Fillatre, L., Vaton, S., Chonavel, T.: Efficient Methods for Traffic Matrix Modeling and On-line Estimation in Large-Scale IP Networks. Proceedings of the ITC 21 (2009)
15. Casas, P., Fillatre, L., Vaton, S.: Robust and Reactive Traffic Engineering for Dynamic Traffic Demands. Proc. NGI'08 (2008). doi:10.1109/NGI.2008.16
16. Casas, P., Fillatre, L., Vaton, S.: Multi Hour Robust Routing and Fast Load Change Detection for Traffic Engineering. Proceedings of the IEEE ICC'08 (2008) doi:10.1109/ICC.2008.1081
17. Ben-Ameur, W.: Between Fully Dynamic Routing and Robust Stable Routing. Proceedings of the DRCN'07 (2007). doi:10.1109/DRCN.2007.4762277
18. Pioro, M., Medhi, D.: Routing, Flow, and Capacity Design in Communication and Computer Networks. Elsevier/Morgan Kaufmann, Amsterdam (2004)
19. Zhang, Y.: Abilene Dataset (2004). http://userweb.cs.utexas.edu/∼yzhang/. Accessed 12 Aug 2010
20. Abilene Obs.: http://www.internet2.edu/observatory. Accessed 12 Aug 2010
21. Coates, M., Hero, A., Nowak, R., Yu, B.: Internet tomography. IEEE Signal Process. Mag. **19**(3), 47–65 (2002)
22. Medina, A., Salamatian, K., Bhattacharyya, S., Diot, C.: Traffic Matrix Estimation: Existing Techniques and New Directions. Proceedings of the SIGCOMM'02 (2002). doi:10.1145/964725.633041
23. Zhang, Y., Roughan, M., Duffield, N., Greenberg, A.: Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Load Measurements. Proceedings of the SIGMETRICS'03 (2003). doi:10.1145/781027.781053
24. Nürnberger, G.: Approximation by Spline Functions. Springer, New York (1989)
25. TOTEM: http://totem.run.montefiore.ulg.ac.be/. Accessed 12 Aug 2010
26. Soule, A., Lakhina, A., Taft, N., Papagiannaki, K., Salamatian, K., Nucci, A., Crovella, M., Diot, C.: Traffic Matrices: Balancing Measurements, Inference and Modeling. Proceedings of the SIGMETRICS'05 (2005). doi:10.1145/1071690.1064259
27. Lakhina, A., Crovella, M., Diot, C.: Characterization of Network-Wide Anomalies in Traffic Flows. Proceedings of the IMC'04 (2004). doi:10.1145/1028788.1028813
28. Nikiforov, I.: A lower bound for the detection/isolation delay in a class of sequential tests. IEEE Trans. Inf. Theor. **49**(11), 3037–3047 (2003)

29. Nikiforov, I.: A simple recursive algorithm for diagnosis of abrupt changes in random signals. IEEE Trans. Inf. Theor. **46**(7), 2740–2746 (2000)
30. Basseville, M., Nikiforov, I.: Detection of Abrupt Changes: Theory and Applications. Prentice Hall, NJ (1993)

Author Biographies

Pedro Casas received an Electrical Engineering degree from the "Universidad de la Republica" (UDELAR) in Montevideo, Uruguay in 2005, and a Ph.D. degree in Computer Sciences from Télécom Bretagne, Brest, France in 2010. He holds a teaching and research Assistant position at UDELAR since 2001, and he is a member of the Electrical Engineering Department at the Engineering Faculty of UDELAR since 2003. He joined the french LAAS-CNRS research laboratory in 2010 as a Postdoctoral Research Fellow. His research interests are related to the statistical characterization and analysis of network traffic, network modeling, anomaly detection, and performance analysis in heterogeneous networks supporting Quality of Service.

Lionel Fillatre received the M.Sc. degree in decision and information engineering and the Ph.D. degree in systems optimization from the University of Technology of Troyes (UTT), France, in 2001 and 2004, respectively. From 2005 to 2007, he worked at Télécom Bretagne, Brest, France, first, as a research engineer in the Computer Science department, then as an Associate Professor in the Signal and Communications department. Since 2007, he is an Associate Professor at the Systems Modeling and Dependability Laboratory, UTT. His current research interests include statistical decision theory, signal and image processing, anomaly detection in traffic flows and information hiding in digital imagery.

Sandrine Vaton obtained the Engineering degree from Télécom Paris in 1994, a M.Sc. degree in applied probabilities from the University of Paris 6 in 1995 and a Ph.D. in signal processing from Télécom Paris in 1998. Since 1999 she is an Associate Professor at Télécom Bretagne in Brest, France. Her main research interests concern statistical network traffic analysis, traffic engineering, Quality of Service, and security in telecommunication networks.

Igor Nikiforov received his M.Sc. degree in automatic control from the Moscow Physical - Technical Institute in 1974, and the Ph.D. in automatic control from the Institute of Control Sciences (USSR Academy of Science), Moscow, in 1981. He joined the University of Technology of Troyes (UTT) in 1995, where he is Professor in the system modeling and dependability laboratory (LM2S), which is a part of the Institute of Charles Delaunay, FRE CNRS 2848. His scientific interests include statistical decision theory, detection/isolation of abrupt changes, fault detection/isolation/reconfiguration, signal, image processing and navigation.