On the Benefits of Time-varying Routing in Realistic Mobile Backhaul Networks

Johannes Lessmann, Yong Cheng, Stefan Pfeiffer and Xavier Costa-Pérez

NEC Laboratories Europe
Heidelberg, Germany

{lessmann, cheng, pfeiffer, costa}@neclab.eu

***Abstract*—Since traffic load in mobile networks typically changes significantly over time, time-varying routing, where routing changes in certain time intervals, seems an obvious solution. Yet multiple researchers have claimed independently in the past, that time-varying routing does not lead to worthwhile benefits given its overhead. In this paper, we study this issue in a systematic way. We argue that previous claims are due to focusing on too narrow optimization metrics and demonstrate the poor expressiveness of such metrics using real mobile backhaul topologies. We propose a novel traffic engineering metric, capacity variation, and on that basis, prove that time-varying routing can in fact have great benefits, particularly in terms of infrastructure CAPEX costs. A quantitative evaluation on the benefits is provided based on real mobile backhaul data from a European operators.**

***Index Terms*—traffic engineering, capacity waste, time-varying routing, mobile backhaul.**

I. Introduction

In a mobile network, traffic from base stations (such as NodeBs, eNBs) towards the core network and vice versa is transported via the so-called mobile backhaul network. The mobile backhaul network typically consists of a wired part (optical fiber) in the aggregation domain and a wireless part (usually microwave radio) in the access domain. There can be hybrid regions with a mix of wired and wireless backhaul links. The core and aggregation domain typically employ ring or mesh topologies. In the access domain, most current deployments are still tree, star or chain topologies, but there is a trend towards ring topologies in the last mile as well, all the more so with the introduction of outdoor small cells.

Given that the backhaul has already started to become the primary capacity bottleneck in mobile networks, efficient resource utilization of the already deployed backhaul links is extremely important. The means of choice for ensuring this is usually traffic engineering (TE). However, in mobile backhaul networks, the traffic load at base stations fluctuates considerably across space and time. Hence, time-varying routing has been proposed [4], where routing changes to adapt to demand. Since traffic in a mobile network typically has a recurring pattern with fixed periodicity (e.g. a weekday) and some "homogeneous" sub-periods (e.g. morning, day, evening, night), time-varying routing typically means that one routing configuration is computed per sub-period, i.e. routing changes are enacted at the transition from one period to the next.

While it seems intuitive that time-varying routing leads to better capacity utilization in the network compared to a static routing (where paths never change regardless of the load distribution), multiple researchers have come to the conclusion in previous works (see Section II) that this is not really the case. Even though it is explicitly admitted that the findings are counter-intuitive, the results indicate that the performance difference between static and time-varying routing is only marginal and not worth the extra effort. The main contribution of this paper is to study this issue more deeply. Our finding is that previous claims are indeed true, if traditional TE metrics are used (as is the case in those works). We therefore introduce another TE metric, capacity variation, which is not only more expressive in terms of the real optimization objective, but also leads to considerable benefits of time-varying over static routing, particularly in terms of network CAPEX costs.

An additional significant value of the work presented in this paper is the fact that our performance evaluation is based on real mobile backhaul topologies from European tier-1 operators on top of an exhaustive set of synthetic topologies. Such data is very important to add considerable realism and credibility to the performance benefits evaluation results. To the authors knowledge there are very few works in the state of the art based on real mobile backhaul topologies.

II. Related Work

The authors of [1] have introduced the idea of exploiting predicted periodic changes in the traffic. They do this by tuning the OSPF weight settings in the network as some way of traffic engineering. The authors have introduced an optimization of this in [2] by proposing a function of load and capacity as metric for the routing algorithm. The problem of using the load as a factor in the metric for the shortest-path algorithm is the introduced instability that can lead to traffic fluctuations. [1] and [2] belong to the category of papers which claim that time-varying routing does not lead to significant benefits.

The authors of [3] have asserted that "... weight setting for OSPF cannot replace MPLS as a traffic engineering tool." MPLS is a prominent technology in mobile backhaul networks, since it combines some of the advantages of packet switched and circuit switched (ATM and TDM) transport networks.

A paper that utilizes MPLS instead of IP routing to reduce the overall costs of operations by more efficient use of bandwidth resources through explicit routing is [4]. The authors of this paper, which is a follow-up of [5], have solved the routing problem on a MPLS network with time-variable traffic demand. The authors have proposed a simple on-line algorithm for optimal selection of the label switched path (LSP) and an offline algorithm that uses an Integer Linear Programming (ILP) formulation. Their proposed approach considers, similar to this paper, multiple traffic matrices for different time slots. Their work is, like most others, based on artificial, non-backhaul topologies. The authors also claim that time-varying routing has only marginal gains compared to static routing.

The work in [6] also proposes time-dependent routing. It proposes a hybrid TE solution with an offline path computation component that tries to globally optimize routing based on varying demand, and an online component for unforeseeable excess traffic. The authors do not study the difference to static routing.

III. Bottleneck Problem in Classical TE

In our analysis of static versus time-varying routing, we will see that previous works' claim of time-varying routing not offering much performance benefit on top of static routing is really due to being based on the classical TE metric of "maximizing the minimum residual link capacity". Hence, as a precursor to our main study, in this section we will make a little foray into the significance of this issue in real backhaul networks. The mentioned TE metric has been around for a long time and is meant to enforce load balancing across the links in the network: there should not be any links with exceptionally high congestion levels compared to other links. If even the most congested link can still offer a maximum of residual capacity, the network is considered to be evenly balanced in terms of link loads.

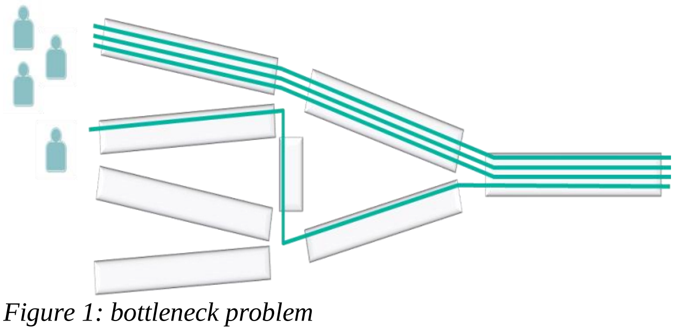

Figure 1: bottleneck problem

It has been argued that such TE metric might not be always ideal. The reason is that it focuses entirely on the characteristics of a single link in the network, namely the most congested one. It does not really take additional links, much less a majority of links, into consideration. The problem with this is depicted in Figure 1, which shows a schematic aggregation network. Since all traffic flows between one of the four edge nodes on the left-hand side of the network and the one edge node on the right-hand side, all traffic must traverse the single rightmost aggregation link. Assuming all links have the same capacity, this implies that the mentioned aggregation link will always be the bottleneck, i.e. the most congested link, no matter how traffic is routed in the rest of the network. Because of that, the TE objective to "maximize the minimum residual link capacity" fails to capture an expressive measure for load balancing in the network as a whole.

The problem that such TE metrics are not expressive for the overall network but will be dominated by only a few bottleneck links has been addressed in the past by proposing alternative TE metrics such as "maximize the *average* residual link capacity". In this paper, we do not claim to be the first ones to ever acknowledge this problem space. Rather, our contribution is twofold: first, we show the extent to which this bottleneck problem is actually present in real-world mobile backhaul networks rather than in synthetic examples. Second, we propose a new TE metric, capacity variation, which is not only more expressive for the overall network, but also has a considerable impact on infrastructure cost. At the same time, we will show that by looking at our new metric, there is indeed a significant performance difference between static and time-varying routing. While we focus on the first contribution in this section, the second one will be dealt with in Section IV.

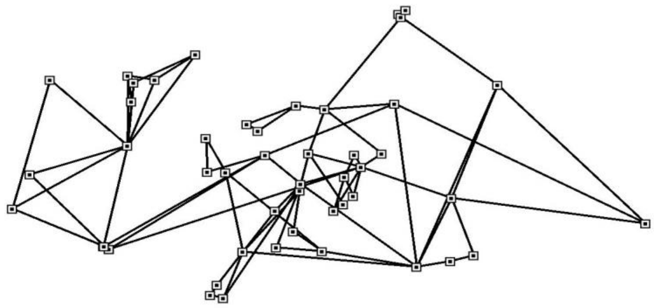

Figure 2: Metropolitan backhaul topology of tier-1 operator

Figure 2 depicts a real microwave aggregation network topology in a large European city. The whole network includes about 470 sites, most of which are connected in a simple hub-and-spoke fashion to the depicted redundant aggregation network which consists of only 48 sites. Within this network, about 20 sites were fully fledged aggregation hubs with collocated access routers. This means that, on average, an access router connects about 24 sites to the RNC, the 3G gateway which is not shown in the figure.

The connectivity of each of the 20 aggregation hubs to the RNC (which is based on fiber) was not known to us at the point of our analysis. Hence, we studied three different variants: 1) we chose two aggregation hubs at random, but far apart from one another (results for this variant will be denoted "2.x"; 2) we chose not two, but three aggregation hubs (results being denoted "3.x"; 3) we chose all aggregation hubs. The chosen aggregation hubs were then connected in our simulations via a virtual link with infinite capacity directly to the RNC. This was to emulate the fact that the fiber network between a set of aggregation hubs and the RNC would in reality never be the capacity bottleneck for the network as a whole. Note that, in the following, we will show results only for variants 1) and 2) due to space limitations. The results for variant 3) are perfectly in line with our presented results and were used by us as an additional verification for the validity of our claims.

We then computed a routing for each of the sites through the aggregation network to the RNC based on traffic information obtained from the operator. The actual traffic demand does not impact the fundamental outcome of our study as long as it remains the same for all simulation variants[1]. The routing we used was based on a Genetic Algorithm (GA) presented in [7], which attempts to find a globally optimum routing for all service demands, while maximizing the minimum residual link capacity, i.e. our TE metric under study. Our goal was to find out to which extent real backhaul topologies exhibit a bottleneck problem, i.e. how much the network throughput is constrained by a few selected links.

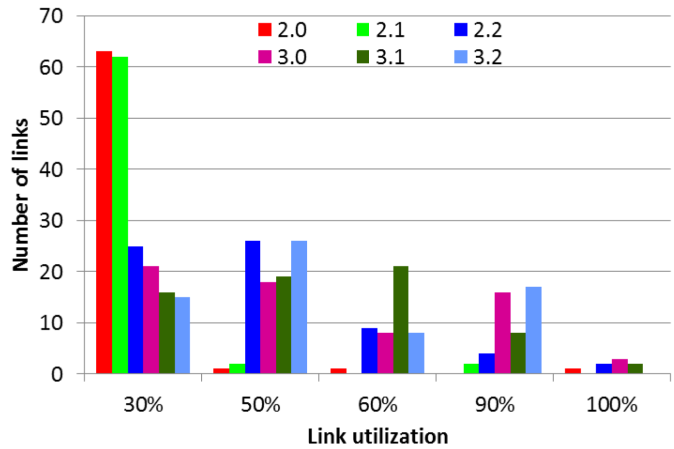

Figure 3: Link utilization histogram for different iterations

After running the routing algorithm in our network, we had a look at the link utilization histogram for all network links which is depicted in Figure 3. As can be seen in the data series corresponding to the “first iteration” (i.e. 2.0 and 3.0 depending on the routing variant), most of the links are only lightly utilized while very few links have an exceptionally high utilization of more than 90% (“100%” on the x-axis actually means “between 90% and 100%”). Clearly, this is a network bottleneck which throttles down the whole network throughput. Apparently, even our fairly sophisticated optimization algorithm could not come up with a better solution in terms of more network throughput or better load balancing. Hence, the identified link seems to indeed be a bottleneck by the way this topology was designed.

To see the impact of fixing this bottleneck problem, we manually increased the capacity of the bottleneck link by a factor of 10, thus ensuring it will not be the limiting entity again. After running the GA again in the modified network, two new bottleneck links emerged. The histogram for the “second iteration” (i.e. 2.1 and 3.1 depending on the routing variant) in Figure 3 still shows a very skewed utilization distribution, yet the total possible network throughput in this case almost doubled compared to the previous run (cf. Figure 4, to be explained later). This shows how severely bottlenecks do limit network throughput in current backhaul deployments.

We fixed the two new bottleneck links again by generously increasing their capacities and performed a third iteration. After this, while there are again bottleneck links in the network, the utilization distribution is much more balanced (cf. values for 2.2 and 3.2 in Figure 3). It becomes clear from this analysis, that networks need to be carefully dimensioned to ensure network throughput is not constrained by only a few bad dimensioned links. Yet the analysis also makes clear that this can in fact be done and existing deployments may need only a few upgrades to a handful of links in order to become far more balanced and capacity-efficient.

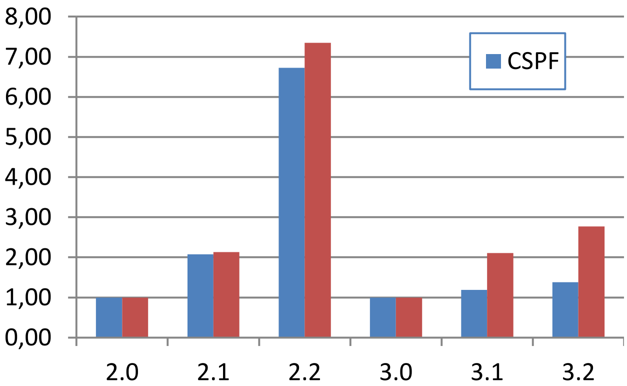

Figure 4: network throughput for variants 2.x and 3.x

Finally, Figure 4 gives an overview of the evolution of achievable network throughput after each bottleneck fixing for our 2 topology variants. The throughput values are normalized with respect to the network throughput which was feasible in the original network (i.e. w.r.t. 2.0 and 3.0). We did this study with Constrained Shortest Path First (CSPF) in addition to the Genetic Algorithm (GA) used so far to verify that the results are not related to the selection of the algorithm. CSPF is a routing strategy employed in many real backhaul networks, which means it does in fact give a realistic impression of potential throughput evolution in our bottleneck study. Clearly, regardless of how well the hubs are connected, the limiting effect of bottleneck links in backhaul networks is remarkable.

An interesting outcome of this study is that the achievable network throughput in a real backhaul network might increase dramatically by just increasing the capacities of very few links in the network. It is also important to provide sufficient connectivity to the RNC, as that will reduce the sensitivity of the network to bottlenecks (the throughput increase for the 3.x series is clearly less than for the 2.x series). This underlines the argument made earlier in this paper, that applying classical TE metrics which focus only on the bottleneck link, is not expressive for the network as whole and might not lead to the desired result of optimum load balancing.

[1] It cannot be disclosed in this paper for confidentiality reasons.

IV. CHANGING THE TE METRIC: CAPACITY VARIATION

In the previous section, we have demonstrated that the bottleneck problem, which is intuitive in theory, is indeed relevant to a significant extent also in real backhaul networks. In this section, we argue that an implication of this problem, namely that certain TE metrics which focus on bottleneck links may have little expressiveness for the overall performance, also extends to the question of whether or not time-varying routing has sufficient advantages over static routing.

In order to study the performance difference of static versus time-varying routing, we computed optimal solutions for both routing variants across an exhaustive set of topologies. This was done to ensure that random effects due to the selection of a particular topology (in which, for instance, time-varying routing happens to perform exceptionally poor) were eliminated. Our approach was therefore as follows. We chose a certain node count *n*. We then generated all possible topologies for an *n*-node network such that the topology is fully connected (i.e. there is no isolated node or cluster of nodes being unconnected to the rest of the topology). For each of these topologies, we generated two Mixed Integer Linear Programs (MILPs): one that models the optimal solution for static routing, the other one that models the optimum solution for time-varying routing for that given topology.

In the *n*-node topology, *n-1* nodes were assumed to be base stations, i.e. generating traffic, while 1 node was assumed to be the gateway, i.e. the traffic sink. The traffic demand was generated quasi-randomly, but in a way that the demand distribution across base stations changed over time. This was done to emulate large-scale user mobility, e.g. from commercial city areas over day to residential areas over night. Once the traffic was generated, it was used for all computations, be it static or time-varying, with a certain node count. To model time variability, we used 3 time slots with each time slot having a distinct traffic distribution (the total aggregate demand hardly changed over time). This is in line with the recurring traffic pattern that mobile backhaul networks typically have. Traffic typically has a fixed periodicity (e.g. a day) and multiple "homogeneous" sub-periods. In our model, we "emulated" the typical periods of day, evening and night.

It is important to understand that static routing can principally be computed in two different ways. One is to start with multiple matrices for multiple time slots capturing demand variations over time and then collapsing these multiple matrices into one computing a single static routing based on that single aggregated matrix. Another way is to compute a static routing (which does not change over time) that is as optimal as possible with respect to all traffic matrices at the same time. In the latter case, time variations in traffic demand are explicitly taken into account for computing the solution, even though the ultimate routing is static. Generally, the latter way of computation leads to far better performance in terms of network throughput than the former way. While operators will typically employ the former strategy, our MILPs are actually based on the latter form of computation. This means that our comparison of static with time-varying routing is very favorable to static routing. In other words, comparing time-varying routing to the former strategy of static routing will yield even greater performance differences than shown in our results.

We will now introduce our MILP for time-varying routing, since this is the more complex case and also involves our new TE metric. Due to space limitations, we cannot introduce our MILP for the static routing case here. It is available upon request to the authors.

minimize $\propto c_{max} + \beta\bar{u}$

s.t.

$$\sum \mathrm{r}_{\mathrm{i,n}}^{\mathrm{i,j,k,s}} = 1 \;\; \forall i,j,k,s,n \tag{1}$$

$$\sum \mathrm{r}_{\mathrm{n,i}}^{\mathrm{i,j,k,s}} = 0 \;\; \forall i,j,k,s,n \tag{2}$$

$$\sum \mathrm{r}_{\mathrm{n,j}}^{\mathrm{i,j,k,s}} = 1 \;\; \forall i,j,k,s,n \tag{3}$$

$$\sum \mathrm{r}_{\mathrm{j,n}}^{\mathrm{i,j,k,s}} = 0 \;\; \forall i,j,k,s,n \tag{4}$$

$$\sum \mathrm{r}_{\mathrm{m,n}}^{\mathrm{i,j,k,s}} = \sum \mathrm{r}_{\mathrm{n,m}}^{\mathrm{i,j,k,s}} \;\; \forall i,j,k,s,m,n \tag{5}$$

$$\sum (\mathrm{d}^{\mathrm{i,j,k,s}}\,\mathrm{r}_{\mathrm{m,n}}^{\mathrm{i,j,k,s}} + \mathrm{d}^{\mathrm{i,j,k,s}}\,\mathrm{r}_{\mathrm{n,m}}^{\mathrm{i,j,k,s}}) = \mathrm{u}^{\mathrm{m,n,s}} \;\; \forall i,j,k,s,m,n \tag{6}$$

$$\mathrm{u}^{\mathrm{m,n,s}} \leq \mathrm{v}^{\mathrm{m,n}} \;\; \forall m,n,s \tag{7}$$

$$\mathrm{v}^{\mathrm{m,n}} \leq \mathrm{c}_{\mathrm{max}}\,\mathrm{C}_{\mathrm{m,n}} \;\; \forall m,n \tag{8}$$

$$0 < \mathrm{c}_{\mathrm{max}} \leq 1 \tag{9}$$

$$\check{\mathrm{u}}^{\mathrm{m,n}} \leq \mathrm{u}^{\mathrm{m,n,s}} \leq \hat{\mathrm{u}}^{\mathrm{m,n}} \;\; \forall m,n,s \tag{10}$$

$$\bar{\mathrm{u}}^{\mathrm{m,n}} = \hat{\mathrm{u}}^{\mathrm{m,n}} - \check{\mathrm{u}}^{\mathrm{m,n}} \;\; \forall m,n \tag{11}$$

$$\bar{\mathrm{u}} = \sum \bar{\mathrm{u}}^{\mathrm{m,n}} \;\; \forall m,n \tag{12}$$

$$\mathrm{r}_{\mathrm{m,n}}^{\mathrm{i,j,k,s}} + \mathrm{r}_{\mathrm{n,m}}^{\mathrm{i,j,k,s}} \leq 1 \;\; \forall i,j,k,s,m,n \tag{13}$$

$$\mathrm{r}_{\mathrm{m,n}}^{\mathrm{i,j,k,s}} \in \{0,1\} \;\; \forall i,j,k,s,m,n \tag{14}$$

$$\sum_{\mathrm{k}=0}^{\mathrm{K}} \mathrm{d}^{\mathrm{i,j,k,s}} = \mathrm{d}^{\mathrm{i,j,s}} \tag{15}$$

$$\mathrm{K} \leq \gamma \tag{16}$$

Equations (1) to (5) are classical network flow constraints. Variable *r* is a binary routing variable (cf. Equation (14)) that is 1 when a traffic demand between node *i* and *j* (as given by the traffic matrix) in time slot *s* is non-zero and routed over link (*m*,*n*), i.e. in that direction. Otherwise, it is zero. Note that we need r to be binary since we do not want to assume a fluid traffic model, i.e. infinite splitting granularity, which is used in many previous works, but is unrealistic. Hence, Equation (1), for instance, says that the sum of routing variables for the k-th fragment of a traffic demand between i and j in time slot s, leaving the source node i of that demand, must amount to 1, i.e. the full demand must leave the source node. Likewise, the full demand must enter the destination j (cf. Equation (3)). Equations (2) and (4) say that demands may not enter its source or leave its destination. Equation (5) is the classical continuity constraint on intermediate nodes for a demand. Clearly, k is bound by the splitting granularity γ, i.e. the maximum number of distinct flow classes that the ingress flow classifier can distinguish (cf. Equation (16)). For the sake of brevity, some constraints in the MILP are left out. Equation (1) is only specified, for example, if the traffic demand $d^{i,j,k,s} > 0$ and $(i, n) \in E$. Similar omissions hold for other equations.

Variable $u^{m,n,s}$ in Equation (6) is the consumed capacity of link (m,n) in time slot s and thus the sum of all demands that are routed over that link in either direction. Equation (7) then defines $v^{m,n}$ as the maximum consumed capacity of (m,n) in any time slot. This in turn is bound by the available physical link capacity $C_{m,n}$ (Equation (8)). Variable c_{max} is the maximum link utilization in the network and will be minimized as part of the objective function. This is another way of maximizing the minimum residual link capacity, i.e. represents a classical TE metric used in previous works on time-varying routing. In order to go beyond this TE metric, our objective function contains another metric, which we introduce now.

The consumed capacity in slot s, $u^{m,n,s}$ is used in Equation (10) to define two other variables, namely the lower bound or minimum consumed capacity across all time slots, $\check{u}^{m,n}$, and the upper bound or maximum consumed capacity across all slots, $\hat{u}^{m,n}$. The difference between the two is defined as $\bar{u}^{m,n}$ in Equation (11), which we call the ***capacity variation*** per link (m,n). Variable $\bar{u}$ (cf. Equation (12)) is then the total capacity variation in the network. Intuitively, the capacity variation of a link (our novel optimization metric) is the difference between the maximum utilization and the minimum utilization of a particular link *across all time slots*. A high value thus means that some load goes over a certain link in time slot t_1, while much less load goes over that same link in another time slot t_2. Clearly, the higher this capacity variation, the higher is the capacity waste. This is because the operator has to invest into a link with sufficient capacity to accommodate the load in slot t_1, while much of the time this capacity is not used at all and a smaller link would have been sufficient. The MILP above therefore minimizes not only the maximum link utilization, but also tries to minimize the capacity variation in the whole network.

Equation (13) is an important constraint which expresses that a demand should not be routed over a link (m,n) and (n,m) at the same time (i.e. in both directions of the same physical link). This constraint is needed, as else a solution will often have exactly that, because adding a demand "twice" to the consumed capacity $u^{m,n,s}$ on a (not so congested) link can reduce the capacity variation $\bar{u}^{m,n}$ on that link (without impacting the maximum link utilization value), thereby optimizing the objective function in an undesired way.

V. Evaluation Results

In order to evaluate the performance difference between static and time-varying routing, we generated all connected topologies given a certain node count, and two MILPs (for both routing variants) for each of those topologies. The number of connected topologies increases exponentially with increasing node count. For 4 nodes, there are 34 connected topologies, for 5 nodes, there are already 728 ones, for 6 nodes, there are 26704 ones, for 7 nodes, there are 1866256 candidates, and so on. Additionally, the number of possible routings for a given topology increases exponentially with increasing node count. Hence, the overall product (#topologies *x* #routings) becomes excessively large for even moderate node counts. This is why we focus our evaluation on 6-node networks (and smaller).

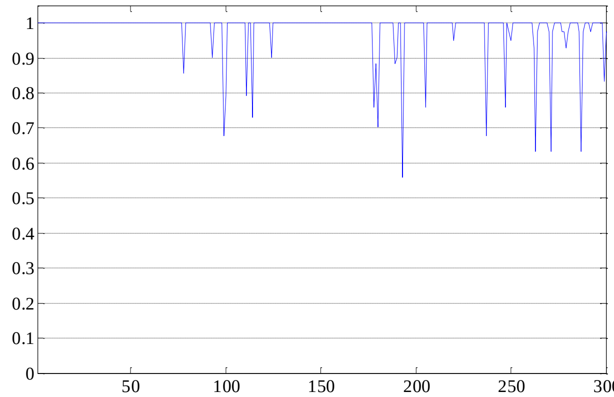

Figure 5: performance ratio for max. link utilization

Figure 5 depicts the performance ratio of time-varying over static routing in terms of the maximum link utilization, i.e. c_{max} in our MILP. According to this metric, better performance translates into smaller values, which means that time-varying routing performs relatively better than static routing if the performance ratio is less than 1. The x-axis enumerates all possible topologies, in this case for a 5-node network. The topologies are sorted according to the average node degree. This means that the left side of the x-axis contains tree topologies (minimum-degree topology that is connected) while topologies get increasingly meshed towards the right side of the x-axis. While there is no path diversity, thus no routing and therefore also no possible performance difference in pure tree topologies, one would expect some benefit of time-varying over static routing in more meshed topologies. As can be seen from Figure 5, however, this is not the case. Even though there are outliers where time-varying routing outperforms, the ratio converges to 1 even in the meshed case. This confirms the findings of previous works which have claimed that time-varying routing is not significantly more effective than static routing in terms of load balancing the traffic.

We now consider the same analysis, but with the new *capacity variation* metric. While in the previous case, we have focused only on minimizing c_{max}, we will now consider the whole objective function as given in Section IV. Figure 6 shows the result. Again, the x-axis enumerates all possible topologies according to average node degree, in this case of a 6-node network. We have performed evaluations for 4-, 5- and 6-node networks, and the results are comparable regardless of the node count.

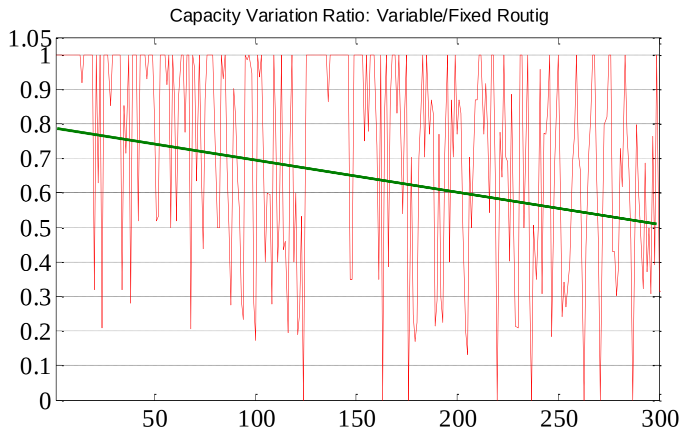

Figure 6: performance ratio for capacity variation

Once we consider the capacity variation metric, we can observe that time-varying routing exhibits a significant performance gain over static routing. To better visualize this, we have included the linear regression line in Figure 6. Clearly, the average gain of time-varying routing trends towards a 50% improvement in highly meshed topologies. This impressive result counters the argument from previous works that time-varying routing does not lead to significant performance gains *in a general sense*. It obviously all depends on the metric one tries to optimize. As we have motivated, capacity variation as a metric has lots of relevance to operators as it minimizes the amount of capacity which an operator needs to "invest" even under unstable traffic conditions.

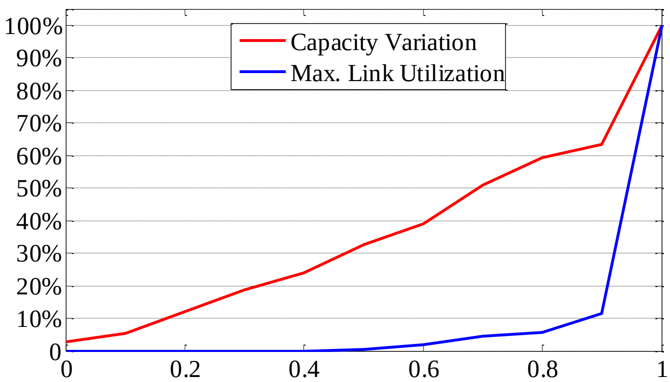

Figure 7: CDF for both TE examined metrics

Figure 7 finally shows the CDF for the link utilization and the capacity variation metric, respectively. The x-axis shows the performance ratio of time-varying over static routing. The numbers were computed analogously to Figures 5 and 6. As can be seen from Figure 7, for only about 10% of all topologies, time-varying routing lead to a maximum link utilization of 90% or less compared to the corresponding value for static routing (i.e. performance improvement of 10% or more in terms of the maximum link utilization metric). With the capacity variation metric, on the other hand, the CDF is much more linear. About 60% of topologies lead to a gain of 10% or more for time-varying routing. About 40% of topologies result in a performance gain of 40% or more.

VI. Conclusion

In this paper, we have studied the performance of time-varying versus static routing. First, we have argued that TE metrics which focus only on bottleneck links can lack expressiveness about the network as a whole. To support this claim (which has been made in the past as well), we have examined the significance of the bottleneck problem in a real-world mobile backhaul topology of a tier-1 European operator. We could show that topologies exhibit bottleneck constraints to a significant extent, and, while better balanced in the end, are still somewhat skewed in terms of link utilization after fixing bottlenecks in the original network several times.

We have then introduced an MILP based on a new TE metric, capacity variation, which minimizes capacity waste and therefore helps reduce infrastructure cost. By means of an exhaustive study with 1000s of topologies, we have analyzed the performance of optimal static and time-varying routing solutions. Our conclusion is that previous claims which have questioned the benefit of time-varying routing have considered a too narrow context. With the capacity variation metric (and arguably other more network-wide metrics as well), time-varying routing leads to considerable performance gains of up to 50% in realistic mobile backhaul topologies.

References

[1] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world", IEEE Journal on Selected Areas in Communications, volume 20, pp. 756-767, May 2002.

[2] B. Fortz, J. Rexford, and M. Thorup, "Traffic Engineering With Traditional IP Routing Protocols", IEEE Communications Magazine, October 2002.

[3] D. Lorenz, A. Ordi, D. Raz, and Y. Shavitt, "How good can IP routing be?", DIMACS Technical Report 2001-17, 2001.

[4] F. Ricciato, U. Monaco, "Routing demands with time-varying bandwidth profiles on a MPLS network", Computer Networks 47(1): 47-61 (2005).

[5] F. Ricciato, S. Salsano, A. Belmonte, M. Listanti, "Off-line Configuration of a MPLS over WDM Network under Time-Varying Offered Traffic", IEEE INFOCOM 2002

[6] Y. Yang, C.-H. Lung, "The Role of Traffic Forecasting in QoS Routing – A Case Study of Time-Dependent Routing", IEEE International Conference on Communications (ICC), 2005

[7] J. Lessmann, "Resource Optimization in Realistic Mobile Backhaul Networks", IEEE ICC 2015