



# Boosting performance for software defined networks from traffic engineering perspective

Mohammed I. Salman<sup>\*</sup>, Bin Wang

Department of Computer Science and Engineering, Wright State University, Dayton, OH, United States of America

## ARTICLE INFO

### Keywords:

Traffic engineering  
Routing schemes  
Software defined networking  
Oblivious routing  
Simulation  
Optimization

## ABSTRACT

Paths selection algorithms and rate adaptation objective functions are usually studied separately. In contrast, this paper evaluates some traffic engineering (TE) systems for software defined networking obtained by combining path selection techniques with average delay and load balancing, the two most popular TE objective functions. Based on TE simulation results, the best TE system suitable for software defined networks is a system where the paths are calculated using an oblivious routing model and its adaptation rate calculated using an average delay objective function. Thus, we propose the RACKe+AD system combining path sets computed using Räcke's oblivious routing and a traffic splitting objective function using average delay. This model outperforms current state-of-the-art models, maximizes throughput, achieves better network resource utilization, and minimizes delay. The proposed system outperformed SMORE and SWAN by 4.2% and 9.6% respectively, achieving 27% better utilization and delivering 34% more traffic with 50% less latency compared with both systems on a GÉANT network.

## 1. Introduction

Centralized traffic engineering (TE) has gained much attention following new software defined networking (SDN) developments. Large technology companies such as Microsoft [1] and Google [2] have shifted to this technology over the last few years.

Some previous studies have deviated from the standard SDN centralization feature to improve scalability and fast adaptation to changing traffic conditions, e.g. Contra [3], HULA [4], MP-HULA [5], and DASH [6] balance load traffic entirely in the data plane to reduce controller overhead. These solutions provide scalable systems with short response time, but degrade performance, with resulting distributed solutions far from optimal [7].

Performance can also be affected by the traffic splitting objective function. Some TE systems balance load over some paths by minimizing maximum link utilization (MLU) [1,8]. However, minimizing MLU does balance load and enhance performance for low traffic and degrades performance significantly during peak hours since it requires additional constraints to satisfy all the demands [9]. Other TE systems use meta-heuristic [10] or heuristic [11] solutions that can provide fast routing convergence, but the solutions are sub-optimal since they may be only local optima. Prior to SDN, several studies considered different objectives [12,13]. To our knowledge, performance impacts from these objectives and path selection strategies have not been properly considered for SDN. Any TE system has two key ingredients: which set

of paths is used for forwarding traffic, and how to split traffic over these selected paths. To the best of our knowledge, no previous study has focused on boosting performance by optimizing combinations of these key ingredients, in contrast, previous work has focused on either path selection algorithms or traffic-splitting objective functions, but not both.

Many studies suggest that a set of shortest paths should be used in TE systems to achieve reliable performance [1,14,15]. Unfortunately, choosing shortest paths may exacerbate congestion for topologies with high link capacity heterogeneity. Oblivious routing<sup>1</sup> strategies offer network demand independent routing schemes, i.e., the routing scheme that is oblivious to the demands [16–19]. Although oblivious routing schemes can be devised with guaranteed congestion ratio, the resulting routing scheme is static and unable to adapt to changing traffic conditions. Several studies have shown that route allocations calculated using an oblivious routing model achieve comparable quality to adaptive solutions [8,20]. Selected paths from this approach are capacity-aware and diverse, which improves not only system performance, but also robustness.

The capacity aware concept not only applies to path selection only, but also to sending rates. For example, the Kleinrock delay objective function [21] minimizes congestion by increasing highly utilized link costs, thus, avoiding highly congested links. The widely used load balancing (LB) objective function [1,8,22–24] minimizes utilization

<sup>\*</sup> Corresponding author.

E-mail addresses: [salman.6@wright.edu](mailto:salman.6@wright.edu) (M.I. Salman), [bin.wang@wright.edu](mailto:bin.wang@wright.edu) (B. Wang).

<sup>1</sup> We use “Oblivious routing” and “Räcke's oblivious routing” interchangeably.

(relative load) for all links, and can also be considered a capacity-aware objective function. The main goal for demand aware objectives is to mitigate proportional increases for all demands [12] by minimizing MLU. However, all source destination (SD) pair demands do not increase at the same rate, and it is not trivial to predict future demands. Thus, sending rates should not only be capacity aware, but also demand aware.

Therefore, we constructed a new simulator, and motivated by SMORE [8] and AD objective functions [23–25] we propose RACKE+AD, a centralized, adaptive, semi-oblivious, demand aware, near optimal TE system with static routes allocated using Räcke's oblivious routing model [16,18,19] and dynamic rate adaptation by approximating the average delay (AD) objective function. RACKE+AD outperformed SWAN [1] and SMORE [8] for throughput, congestion, and latency evaluated on GÉANT and ATT topologies.

**Contributions.** Critical contributions from the current paper are as follows:

- (1) We present a routing scheme that outperforms current state-of-the-art techniques.
- (2) We introduce RACKE+AD, a new efficient TE simulator that can test many routing schemes simultaneously. RACKE+AD is optimized for testing different route selection algorithm and objective function combinations and can be easily extended to test future TE systems.
- (3) We demonstrate that a TE system with static routes and adaptive traffic splitting offers many benefits, including performance, throughput, and resource utilization.

## 2. System model

All TE systems comprise two phases: identifying a set of paths to be used to forward traffic (path selection), and identifying splitting ratios to distribute traffic over these paths (rate adaptation). Generally, routes selected in the path selection phase are static, i.e., selected once and only recalculated when the network topology changes. Path selection is usually offline because updating end-to-end paths may take hundreds of seconds for wide area networks. In contrast, the rate adaptation phase must update path weights regularly due to frequent demand changes. However, the time required to update path weights is considerably less than the time required to update paths in the network. Among many techniques of paths selection algorithms and rate adaptation objective functions, the aim of this research is to find the best combination of these phases to enhance network performance.

**Path and Rate Adaptation Properties:** Intuitively, independently chosen paths may not provide better performance than dependently chosen paths. However, SMORE showed that path selection has considerable effect on performance [8]. Selected paths should be low stretch to minimize latency and naturally load balanced to provide better performance. Low stretch motivated us to compare SMORE performance and latency against k-shortest paths (KSP) approaches. SMORE is naturally load balanced since route computation in Räcke's oblivious routing model is not independent and incorporates some randomness, i.e., the obtained route set may not be the same if we were to run the model again. Thus, we expect different performance for each run. On the other hand, KSP selected paths are not capacity aware due to the natural load balancing. Performance can be further boosted if we use the same concept for splitting traffic over the selected paths, and we expect best performance may be achieved using phases, path selection, and rate adaptation.

### 2.1. Rate adaptation models

#### 2.1.1. Load balance

The load balance (LB) objective is also known as minimizing MLU, Wozencraft objective [26], or minimizing congestion, where LB minimizes the load on the most congested link. Thus, the LB problem can




Fig. 1. Piecewise linear approximation of the delay function.

be expressed as [24]

$$\min F(x) = r \quad (1)$$

$$\text{s.t.} \quad \sum_{p \in P_d} x_{dp} = h_d, \quad d \in D \quad (1a)$$

$$\sum_{d \in D} \sum_{p \in P_d} \delta_{dpl} x_{dp} \leq c_l r, \quad l \in L \quad (1b)$$

where:  $x_{dp}$  is the flow on path  $p$  for demand  $d$ ;  $h_d$  is the volume for demand  $d$ ;  $c_l$  is the capacity for link  $l$ ;  $P_d$  is the number of candidate paths for demand  $d$ ;  $\delta_{dpl} = 0, 1$  is a link-path indicator, with  $\delta_{dpl} = 1$  if path  $p$  for demand  $d$  uses link  $l$ , and 0 otherwise.

Two constraints are applied. The demand constraint (1a) ensures that all demands are satisfied over some paths. The capacity constraint (1b) ensures that load does not exceed the link capacity where  $r \leq 1$ , after solving (1). The linear program formulation above is the final form of the problem whereas the original problem is non-linear. The reader is referred to Chapter 4 of [24] for details on how the problem can be converted to the current form.

#### 2.1.2. Average delay

For this objective function, delay for any network link can be modeled as  $y/(c-y)$ , as shown in (Fig. 1, solid line). Similar to the LB objective, the original AD problem is non-linear and cannot be formulated directly as a linear program. Thus, the delay function is a piecewise linear approximation (2) (Fig. 1, dotted line)

$$g(z) = \begin{cases} (3/2)z & \text{for } 1 \leq z < 1/3 \\ (9/2)z - 1 & \text{for } 1/3 \leq z < 2/3 \\ 15z - 8 & \text{for } 2/3 \leq z < 4/5 \\ 50z - 36 & \text{for } 4/5 \leq z < 9/10 \\ 200z - 171 & \text{for } 9/10 \leq z < 19/20 \\ 4000z - 3781 & \text{for } z \geq 19/20 \end{cases} \quad (2)$$

The linear program for this AD problem is

$$\min F = \sum_{l=1}^L \frac{r_l}{c_l} \quad (3)$$

$$\text{s.t.} \quad \sum_{p=1}^{P_d} x_{dp} = h_d, \quad d = 1, 2, \dots, D \quad (3a)$$

$$\sum_{d=1}^D \sum_{p=1}^{P_d} \delta_{dpl} x_{dp} = y_l, \quad l = 1, 2, \dots, L \quad (3b)$$

$$r_l \geq \frac{3}{2}y_l, \quad l = 1, 2, \dots, L \quad (3c)$$

$$r_l \geq \frac{9}{2}y_l - c_l, \quad l = 1, 2, \dots, L \quad (3d)$$

$$r_l \geq 15y_l - 8c_l, \quad l = 1, 2, \dots, L \quad (3e)$$

$$r_l \geq 50y_l - 36c_l, \quad l = 1, 2, \dots, L \quad (3f)$$

$$r_l \geq 200y_l - 171c_l, \quad l = 1, 2, \dots, L \quad (3g)$$

$$r_l \geq 4000y_l - 3781c_l, \quad l = 1, 2, \dots, L \quad (3h)$$

$$x_{dp} \geq 0, \quad p = 1, 2, \dots, P_k, d = 1, 2, \dots, D \quad (3i)$$

$$y_l \geq 0, \quad l = 1, 2, \dots, L \quad (3j)$$

which is considerably more accurate [24] than the Fortz et al. [27] approximation.

## 2.2. Paths selection algorithms

### 2.2.1. Räcke's oblivious routing model

Räcke's oblivious routing model iteratively computes a distribution over randomized routing trees using an approximation algorithm. Link weights are adjusted for each iteration based on how much the link has been utilized in previous routing tree sets. A routing tree has leaves corresponding to nodes in the original topology. Thus, a path can be obtained between nodes  $u$  and  $v$  in the original graph by finding corresponding leaves for  $u$  and  $v$  in the routing tree.

However, paths for Räcke's oblivious routing model are computed without considering demands, thus, they do not overfit to a specific scenario [8]. Similar to SMORE, we also adopt the simple mechanism used to impose the number of paths for each SD node pair. We use 4 paths for each SD pair of nodes that have the highest weights.

### 2.2.2. $k$ -shortest paths

The proposed KSP algorithm is based on Yen's algorithm, the most commonly used algorithm for TE. KSP is a generalization of the shortest path routing problem. The algorithm returns loopless  $k$  shortest paths ordered from shortest to longest. We use four paths for each SD pair, i.e.,  $k = 4$ .

## 3. Simulator framework

We built a simulator to model and test different TE scenarios, with particular attention to efficiency, simplicity, and extensibility. Although many network simulators have been proposed previously [28–31], they are generally not optimized for modeling TE approaches and/or do not provide ease of use or extensibility. The proposed simulator was built in Python and can test many TE models in parallel while recording statistics in the background. We use Gurobi optimization [32] to solve the linear programming problems, by integrating it with Python. The framework, data and Räcke's oblivious routing model implementation are all available online.<sup>2</sup>

Simulator inputs, (e.g. topology, demands, path selection algorithms, objective functions, etc.) are all specified in a Python script or configuration file. The simulation produces visualized throughput graphs for each TE system. The graphs are updated periodically as throughput data becomes available. Three time-series metrics for each TE system are recorded in the background during simulation: overall throughput, congestion per link, and latency per path. Topology and traffic matrices are provided as input files, where the user provides the location to these files in the configuration file. If the locations are unavailable, random topology and traffic matrices will be generated according to provided parameters, including number of nodes  $N$ , number of links  $L$ , and traffic distribution matrix.

**Table 1**  
Implemented TE algorithms.

TE system	Description
KSP+LB	$k$ -Shortest Paths (KSP) for paths, LB for weights
KSP+AD	$k$ -Shortest Paths (KSP) for paths, AD for weights
RACKE+LB	Räcke's oblivious routing for paths, LB for weights
RACKE+AD	Räcke's oblivious routing for paths, AD for weights
OPTIMAL(LB) <sup>a</sup>	All paths, LB for weights
OPTIMAL(AD) <sup>b</sup>	All paths, AD for weights

<sup>a</sup>The best load balance is achieved with this system.

<sup>b</sup>The best average delay is achieved with this system.

## 4. Simulation setup

### 4.1. Evaluating routing scheme quality

We evaluate TE systems based on congestion, throughput, and delay. Congestion reflects how a TE system utilizes network resources, and we mostly care about congestion when traffic demand exceeds link capacity. Thus, avoiding congestion can be considered as preserving as much residual capacity as possible, which is important for unexpected traffic surges that could cause bottlenecks. Congestion has negative impact on delay due to queuing. We measure path delay by summing queuing delay for each link along that path,  $l/(c-l)$ , where  $l$  is the absolute link load and  $c$  is the link capacity. Throughput is the proportion of total demand that is successfully delivered to the destinations.

### 4.2. Simulation settings

**Path selection algorithms.** We use three approaches for path selection (i) paths selected using Räcke's oblivious routing model, (ii) paths selected using KSP algorithm, and (iii) select all available simple paths. We refer to these RACKE, KSP, and OPTIMAL, respectively.

**Rate adaptation objective functions.** We use two objective functions for rate adaptation: AD and LB. We refer to a routing scheme with paths selected using KSP and rate adaptation using LB objective function as KSP+LB. Similarly, models where the routing scheme selects all available paths and rate adaptation uses AD is referred to as OPTIMAL (AD), etc. The RACKE+LB routing scheme parallels that used in SMORE [8], and KSP+LB is an approximation to the SWAN scheme [1]. Table 1 shows the TE systems used in our experiment.

**Path budget.** Similar to SMORE and SWAN, and to ensure a fair comparison, we use 4 paths to evaluate any routing scheme. If the Räcke's oblivious routing model produces a routing scheme with SD pairs that has more than 4 paths, we use the 4 highest weight paths, similar to SMORE.

**Traffic matrix generation.** We use the gravity model to generate the traffic matrix (TM) [8,17]. The gravity model approximates real-world TMs for a production network [33]. TMs are deduced based on incoming/outgoing flow for each forwarding device. Since that information is not available, we use a capacity based heuristic rather than incoming/outgoing flow information [17].

**Topologies.** We evaluate many TE systems for ATT and GÉANT<sup>3</sup> production topologies. The GÉANT network (European academic network) contains 38 nodes and 104 directed links with heterogeneous capacities. Fig. 2 shows the link capacity distribution for this network. Different TE systems may behave differently depending on link capacity distributions. Shortest-path TE systems may introduce a bottleneck in heterogeneous link capacities as many SD pairs compete for the same resources.

<sup>2</sup> <https://github.com/MohammedSalman/TE-SIMULATOR>.

<sup>3</sup> Dataset available at: <http://www.topology-zoo.org/dataset.html>.

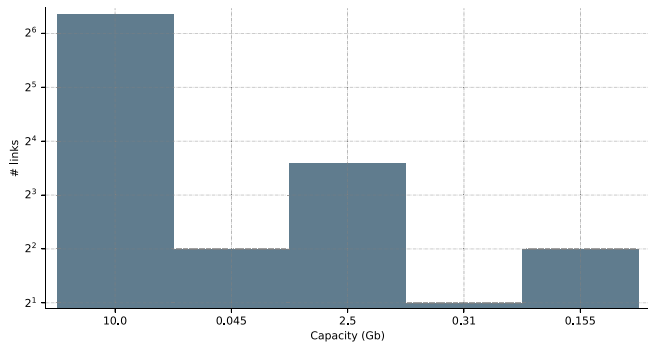


Fig. 2. Capacity distribution for GÉANT network (log scaled).

## 5. Results

We evaluated multiple routing schemes using criteria focused on:

- how each TE system performs regarding throughput and congestion, and
- SMORE and KSP TE system impacts on latency.

### 5.1. Throughput


Performance for many TE systems were evaluated on GÉANT and ATT networks with path budget = 4 for a fair comparison with SMORE. Figs. 3a and 3b show throughput and corresponding throughput distribution for GÉANT network, respectively. Rate adaptation using AD objective function significantly increases throughput, achieving 4.2% and 9.6% improvement over SMORE and KSP+LB, respectively, which confirms path selection effectiveness using Räcke’s oblivious routing algorithm.

Similar to GÉANT topology, a higher throughput was achieved for ATT topology using the AD adaptation rate objective function. KSP had slightly better throughput than Räcke’s oblivious routing path selection algorithm (Figs. 4a and 4b).

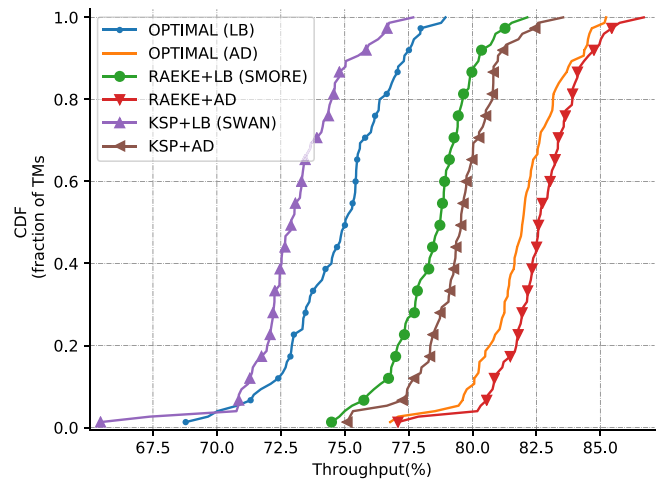
Räcke’s oblivious routing model with LB adaptation rate performed 1.14% better than KSP on average. This may confirm that AD favors shortest paths when all links have the same capacity. However, there is no guarantee that SMORE will always outperform (or underperform) KSP under the same conditions due to oblivious routing scheme randomness. Fig. 5 shows throughput distributions for KSP+AD with a different Räcke’s oblivious routing TE systems obtained by repeatedly calculating the oblivious routing scheme. Output from KSP+AD remained constant since KSP+AD is deterministic. Räcke’s oblivious routing scheme outperformed KSP for 5 runs and underperformed for 1 run. Thus, there is a worst case scenario where KSP may perform better than SMORE. The best run had 2.29% higher throughput than KSP+AD. Therefore, a network operator may choose to run Räcke’s scheme several times and choose the best outcome.

### 5.2. Congestion

Figs. 6a and 6b show network congestion for GÉANT topology using AD and LB. The AD objective function scheduled link loading differently from LB. Fig. 6a shows the maximum congested link over time. All TE systems scheduled link loads that exceeded specific link capacities since we deliberately fed the system with high volume demands to investigate TE system performance well under stressed conditions. AD (Fig. 6a) seems to have higher MLU whereas (Fig. 6b) shows that the AD objective utilizes link loads much better than LB. TE systems with LB caused a bottleneck for more than 40% of links whereas TE systems with AD objective caused a bottleneck for 13% of links. This low congestion ratio for AD is the main reason for the higher throughput (Fig. 3).



(a) Throughput



(b) Throughput distribution

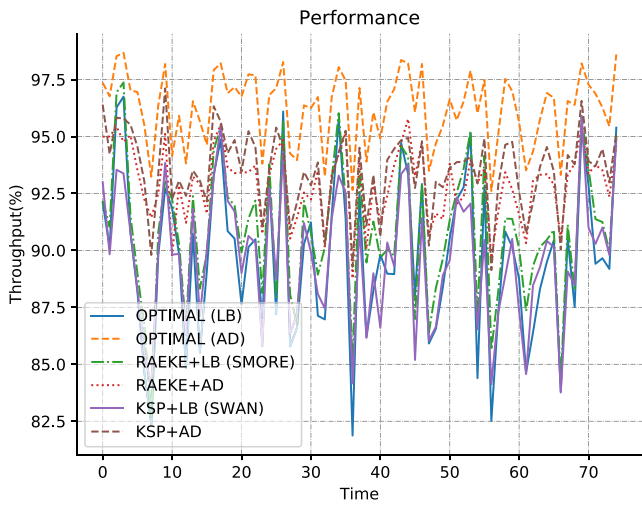
Fig. 3. Throughput for GÉANT topology.

The LB objective always distributes traffic perfectly across the available routes, in the sense that all paths are used and all nodes send and receive traffic with quite similar link utilization (relative load) for all links. Thus, all links might be over-utilized under high demands when the system is not feasible. On the other hand, AD deals more with delay and throughput, but generates worse MLU than from LB. However, MLU is not a true network metric as it only considers congestion for a single link rather than the whole network. Thus, congestion distribution seems like a more reasonable metric, and we only measured MLU to make that point since it is heavily used in the literature.

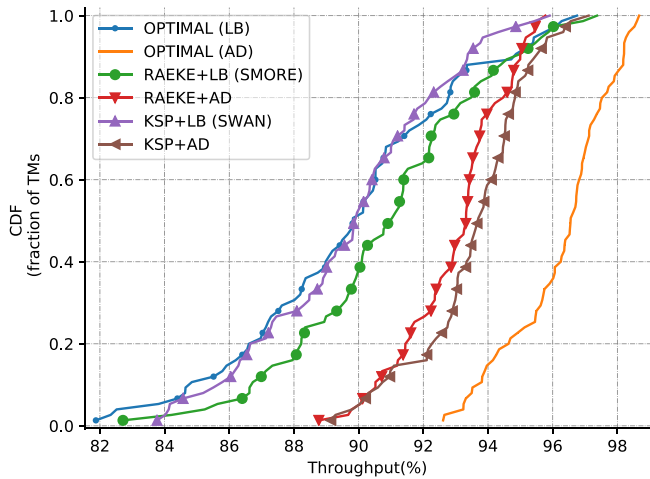
Thus, two factors contributed to better throughput and less congestion: routes selected using Räcke’s oblivious routing algorithm, and using the AD objective. Similar results were obtained for ATT topology (Figs. 7a and 7b).

### 5.3. Latency

Fig. 8 shows link delay distribution with respect to traffic delivered within that delay for GÉANT and ATT topologies. Latency for each path was computed by summing the link delays to obtain the path delay. Including AD selection outperforms LB, achieving significantly lower latency. Fig. 8a shows that LB and AD TE systems different considerably



(a) Throughput



(b) Throughput distribution

Fig. 4. Throughput on ATT topology.

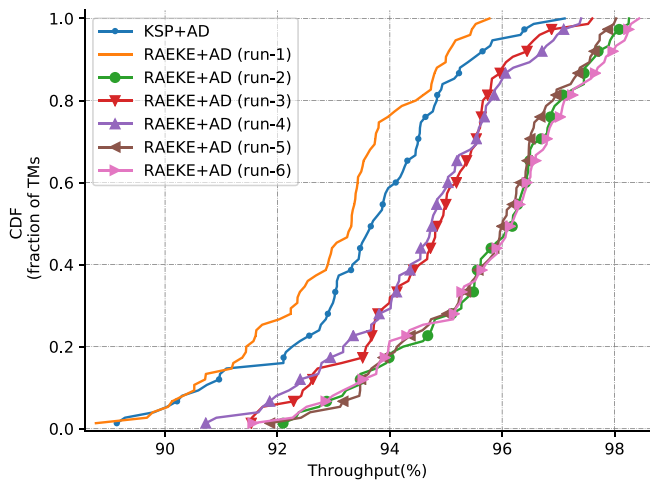




Fig. 5. Throughput distribution for ATT topology for 1 KSP and 6 RAEKE schemes.



(a) Max link congestion, GÉANT topology



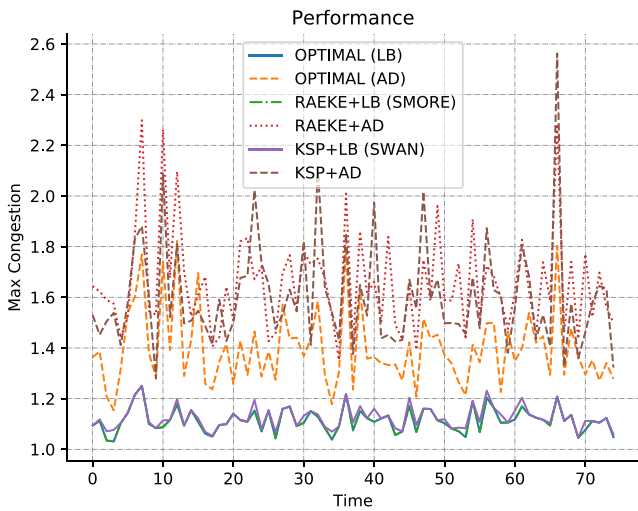
(b) CDF of link congestions, GÉANT topology

Fig. 6. Max link congestion and links' congestion distribution on GÉANT topology.


for GÉANT topology. TE systems with AD objective initially deliver approximately 34% traffic more than those with LB objective, which also has latency 50% lower latency than TE systems with AD. RAEKE+AD routing delivered slightly more traffic than OPTIMAL(AD) since OPTIMAL(AD) goal is to reduce total delay rather than throughput. Fig. 8b shows that routing schemes with AD also delivered more traffic than those with LB for ATT topology. However, the gap between the two groups is somewhat smaller than for GÉANT topologies (Fig. 8(a)) because ATT network links are heterogeneous, hence smaller performance differences between individual links.

## 6. Related work

The classic approach for TE problems is to solve them as a linear program (LP) [24,26], referred to as a multi-commodity flow problem, where the objective function usually minimizes MLU. The approximation of AD objective function is not as widely as used. However, this classical approach does not consider decoupling TE system phases because all available paths are provided as inputs. Choosing all available paths has two limitations: more paths means more decision variables in



(a) Max link congestion, ATT topology



(b) CDF of link congestions, ATT topology


Fig. 7. Max link congestion and links' congestion distribution on ATT topology.

the LP, and forwarding devices, such routers and switches, have limited TCAM memory, hence fewer number paths is always preferable to keep the routing table as small as possible.


The conventional approach adjusts link weights to find a good routing scheme that can increase throughput or minimize congestion in the network [27,34]. However, OSPF can never reach optimal because it uses the equal cost multi-path approach that splits traffic evenly among available shortest paths without rate adaptation. Furthermore, optimizing link weights is an NP-hard problem.

Potentially centralized TE approaches recently became viable due to software-defined networking (SDN) developments, that clearly decouple the two TE phases. SWAN [1] distributes traffic over a set of k-shortest paths using an LP that reserves a small amount of scratch capacity on links to apply updates in a congestion-free manner. SOL [22] uses a greedy approach to randomly select paths with the promise that this random selection will help load balancing traffic across the network. This latter approach is somewhat similar to valiant load balancing [35] but can lead to unnecessarily long paths and consequently increased latency.

Oblivious routing [16–18] has also been proposed to find a routing scheme that performs well under all possible demands. The Räcke



(a) GÉANT topology.



(b) ATT topology.

Fig. 8. Latency distribution.

oblivious routing model [16] guarantees a congestion rate that is never worse than  $O(\log n)$  of optimal, where  $n$  is the number of nodes in the graph. However, despite the guaranteed congestion ratio, this approach cannot outperform systems like SWAN since it considers all possible traffic demands. On the other hand, the oblivious routing approach has inspired several studies (including the current study) to investigate a careful path selection approach. SMORE [8] was inspired by Räcke's oblivious routing model to carefully select paths that increase TE system performance and robustness. Paths selected this way have low stretch, which is important to decrease latency, and are capacity aware, which is important for load balancing. The proposed approach in this paper suggests that careful route selection is not sufficient performance enhancement to reach the expected maximum performance. However, a different objective function from the commonly employed LB could further enhance performance. Hence we were inspired to compare LB and AD objective function performance, and subsequently propose the RAEKE+AD TE system using oblivious routing for path selection with AD to achieve better link delay and network performance.

## 7. Discussion

This section discusses the reason behind the high gap in performance and delay between LB and AD objective functions and one

potential limitation for this work. The LB objective function tends to make the relative load the same for all links when all SD pairs are sending and receiving traffic. This can enhance performance to some extent but causes bottlenecks between some SD pairs under stressed conditions and unpredicted demands, with consequential congestion loss. On the other hand, the AD objective function increases the cost for highly utilized links to avoid utilizing them if other less heavily utilized links are available. Thus, AD is more demand aware than LB and hence offers better contribution to performance. However, solving LP for LB is much faster than for AD, particularly for larger networks due to the increased number of constraints and decision variables.

## 8. Conclusion

Although a few TE systems have been optimized previously using different path selection algorithms, few studies have investigated performance enhancement by testing many objective functions for splitting traffic. These phases have only been studied in isolation previously, with no prior studies testing all possible combinations to find a routing scheme with the best available performance.

This paper proposed RACKE+AD TE system and validated its performance advantages by testing many possible combinations. RACKE+AD selects routes using Räcké's oblivious routing model and adapt traffic using the average delay objective function. Although the intuitive AD goal is to minimize network delay, it also provides surprisingly better throughput than minimizing MLU (commonly known as load balancing).

Simulations confirmed the proposed RACKE+AD system outperformed state-of-the-art routing TE systems in terms of throughput, congestion, and delay. We discussed a caveat when running Räcké's oblivious routing model, where k-shortest paths may give better performance due to randomness in oblivious routing, and also discussed the importance of excluding the maximum congestion metric when evaluating TE systems, particularly system that split traffic not based on the LB objective function.

## CRedit authorship contribution statement

**Mohammed I. Salman:** Conceptualization, Methodology, Software, Data curation, Writing - original draft, Visualization. **Bin Wang:** Validation, Formal analysis, Writing - review & editing, Project administration, Supervision, Resources.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments and suggestions. We also would like to thank Praveen Kumar from Cornell University for addressing all the questions we had regarding the SMORE traffic engineering system.

## References

- [1] Chi-Yao Hong, Srikanth Kandula, Ratul Mahajan, Ming Zhang, Vijay Gill, Mohan Nanduri, Roger Wattenhofer, Achieving high utilization with software-driven wan, in: Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, in: SIGCOMM '13, Association for Computing Machinery, Hong Kong, China, ISBN: 9781450320566, 2013, pp. 15–26, <http://dx.doi.org/10.1145/2486001.2486012>.
- [2] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, Jon Zolla, Urs Hölzle, Stephen Stuart, Amin Vahdat, B4: Experience with a globally-deployed software defined wan, SIGCOMM Comput. Commun. Rev. (ISSN: 0146-4833) 43 (4) (2013) 3–14, <http://dx.doi.org/10.1145/2534169.2486019>.
- [3] Kuo-Feng Hsu, Ryan Beckett, Ang Chen, Jennifer Rexford, David Walker, Contra: A programmable system for performance-aware routing, in: 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), USENIX Association, Santa Clara, CA, ISBN: 978-1-939133-13-7, 2020, pp. 701–721, URL: <https://www.usenix.org/conference/nsdi20/presentation/hsu>.
- [4] Naga Katta, Mukesh Hira, Changhoon Kim, Anirudh Sivaraman, Jennifer Rexford, Hula: Scalable load balancing using programmable data planes, in: Proceedings of the Symposium on SDN Research, in: SOSR '16, Association for Computing Machinery, Santa Clara, CA, USA, ISBN: 9781450342117, 2016, <http://dx.doi.org/10.1145/2890955.2890968>.
- [5] Cristian Hernandez Benet, Andreas J. Kassler, Theophilus Benson, Gergely Pongracz, Mp-hula: Multipath transport aware load balancing using programmable data planes, in: Proceedings of the 2018 Morning Workshop on in-Network Computing, in: NetCompute '18, Association for Computing Machinery, Budapest, Hungary, ISBN: 9781450359085, 2018, pp. 7–13, <http://dx.doi.org/10.1145/3229591.3229596>.
- [6] Kuo-Feng Hsu, Praveen Tammana, Ryan Beckett, Ang Chen, Jennifer Rexford, David Walker, Adaptive weighted traffic splitting in programmable data planes, in: Proceedings of the Symposium on SDN Research, in: SOSR '20, Association for Computing Machinery, San Jose, CA, USA, ISBN: 9781450371018, 2020, pp. 103–109, <http://dx.doi.org/10.1145/3373360.3380841>.
- [7] Mohammad Alizadeh, Tom Edsall, Sarang Dharmapurikar, Ramanand Vaidyanathan, Kevin Chu, Andy Fingerhut, Vinh The Lam, Francis Matus, Rong Pan, Navindra Yadav, George Varghese, Conga: Distributed congestion-aware load balancing for datacenters, in: Proceedings of the 2014 ACM Conference on SIGCOMM, in: SIGCOMM '14, Association for Computing Machinery, Chicago, Illinois, USA, ISBN: 9781450328364, 2014, pp. 503–514, <http://dx.doi.org/10.1145/2619239.2626316>.
- [8] Praveen Kumar, Yang Yuan, Chris Yu, Nate Foster, Robert Kleinberg, Petr Lapukhov, Chiu Lin Lim, Robert Soulé, Semi-oblivious traffic engineering: The road not taken, in: Proceedings of the 15th USENIX Conference on Networked Systems Design and Implementation, in: NSDI'18, USENIX Association, Renton, WA, USA, ISBN: 9781931971430, 2018, pp. 157–170.
- [9] C. Zhang, S. Zhang, Y. Wang, W. Li, B. Jin, R.K.P. Mok, Q. Li, H. Xu, Scalable traffic engineering for higher throughput in heavily-loaded software defined networks, in: NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium, 2020, pp. 1–7, <http://dx.doi.org/10.1109/NOMS47738.2020.9110259>.
- [10] M.M. Tajiki, B. Akbari, M. Shojafar, S.H. Ghasemi, M.L. Barzandeh, N. Mokari, L. Chiaraviglio, M. Zink, Cect: computationally efficient congestion-avoidance and traffic engineering in software-defined cloud data centers, Cluster Comput. (ISSN: 1573-7543) 21 (4) (2018) 1881–1897, <http://dx.doi.org/10.1007/s10586-018-2815-6>.
- [11] W. Quan, N. Cheng, M. Qin, H. Zhang, H.A. Chan, X. Shen, Adaptive transmission control for software defined vehicular networks, IEEE Wirel. Commun. Lett. 8 (3) (2019) 653–656, <http://dx.doi.org/10.1109/LWC.2018.2879514>.
- [12] Eric Gourdin, Olivier Klopfenstein, Comparison of different qos-oriented objectives for multicommodity flow routing optimization, in: Proceedings of the International Conference on Telecommunications (ICT), 2006.
- [13] Simon Balon, Fabian Skivée, Guy Leduc, How well do traffic engineering objective functions meet te requirements?, in: Fernando Boavida, Thomas Plagemann, Burkhard Stiller, Cedric Westphal, Edmundo Monteiro (Eds.), NETWORKING 2006. Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN: 978-3-540-34193-2, 2006, pp. 75–86.
- [14] B. Fortz, J. Rexford, M. Thorup, Traffic engineering with traditional IP routing protocols, IEEE Commun. Mag. 40 (10) (2002) 118–124.
- [15] Zhanyou Ye, Shi Hong Marcus Wu, Themistoklis Prodromakis, Computing shortest paths in 2d and 3d memristive networks, in: Andrew Adamatzky, Leon Chua (Eds.), Memristor Networks, Springer International Publishing, Cham, ISBN: 978-3-319-02630-5, 2014, pp. 537–552, [http://dx.doi.org/10.1007/978-3-319-02630-5\\_24](http://dx.doi.org/10.1007/978-3-319-02630-5_24).
- [16] H. Racke, Minimizing congestion in general networks, in: The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings, 2002, pp. 43–52.
- [17] D. Applegate, E. Cohen, Making routing robust to changing traffic demands: Algorithms and evaluation, IEEE/ACM Trans. Netw. (ISSN: 1558-2566) 14 (6) (2006) 1193–1206, <http://dx.doi.org/10.1109/TNET.2006.886296>.
- [18] Harald Räcké, Optimal hierarchical decompositions for congestion minimization in networks, in: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, in: STOC '08, Association for Computing Machinery, Victoria, British Columbia, Canada, ISBN: 9781605580470, 2008, pp. 255–264, <http://dx.doi.org/10.1145/1374376.1374415>.
- [19] Philipp Zerner, Harald Räcké, Compact oblivious routing in weighted graphs, 2020, arXiv:2007.02427. [cs.NI].
- [20] M. Chiesa, G. Rétvári, M. Schapira, Oblivious routing in ip networks, IEEE/ACM Trans. Netw. 26 (3) (2018) 1292–1305, <http://dx.doi.org/10.1109/TNET.2018.2832020>.
- [21] Leonard Kleinrock, Communication Nets; Stochastic Message Flow and Delay, Dover Publications, Inc., USA, ISBN: 0486611051, 1972.

- [22] Victor Heorhiadi, Michael K. Reiter, Vyas Sekar, Simplifying software-defined network optimization using sol, in: Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation, in: NSDI'16, USENIX Association, Santa Clara, CA, ISBN: 9781931971294, 2016, pp. 223–237.
- [23] X. Liu, S. Mohanraj, M. Pióro, D. Medhi, Multipath routing from a traffic engineering perspective: How beneficial is it?, in: 2014 IEEE 22nd International Conference on Network Protocols, (ISSN: 1092-1648) 2014, pp. 143–154, <http://dx.doi.org/10.1109/ICNP.2014.34>.
- [24] Deep Medhi, Karthik Ramasamy, Chapter 4 - network flow models, in: Deep Medhi, Karthik Ramasamy (Eds.), Network Routing (Second Edition), second ed., in: The Morgan Kaufmann Series in Networking, Morgan Kaufmann, Boston, ISBN: 978-0-12-800737-2, 2018, pp. 114–157, <http://dx.doi.org/10.1016/B978-0-12-800737-2.00005-3>, URL: <http://www.sciencedirect.com/science/article/pii/B9780128007372000053>.
- [25] B. Fortz, L. Gouveia, M. Joyce-Moniz, On the convex piecewise linear unsplittable multicommodity flow problem, in: 2016 12th International Conference on the Design of Reliable Communication Networks (DRCN), 2016, pp. 9–13, <http://dx.doi.org/10.1109/DRCN.2016.7470829>.
- [26] R.G. Gallager, A. Segall, J.M. Wozencraft, Data network reliability, Massachusetts Inst. of Tech. Report, 1976.
- [27] B. Fortz, M. Thorup, Internet traffic engineering by optimizing ospf weights, in: Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064), Vol. 2, 2000, pp. 519–528, <http://dx.doi.org/10.1109/INFCOM.2000.832225>.
- [28] George F. Riley, Thomas R. Henderson, The ns-3 network simulator, in: Klaus Wehrle, Mesut Güneş, James Gross (Eds.), Modeling and Tools for Network Simulation, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN: 978-3-642-12331-3, 2010, pp. 15–34, [http://dx.doi.org/10.1007/978-3-642-12331-3\\_2](http://dx.doi.org/10.1007/978-3-642-12331-3_2).
- [29] Opend Projects, <http://opnetprojects.com/>, URL <http://opnetprojects.com/>.
- [30] Praveen Kumar, Chris Yu, Yang Yuan, Nate Foster, Robert Kleinberg, Robert Soulé, Yates: Rapid prototyping for traffic engineering systems, in: Proceedings of the Symposium on SDN Research, in: SOSR '18, Association for Computing Machinery, Los Angeles, CA, USA, ISBN: 9781450356640, 2018, <http://dx.doi.org/10.1145/3185467.3185498>.
- [31] Steven Gay, Pierre Schaus, Stefano Vissicchio, REPETITA: repeatable experiments for performance evaluation of traffic-engineering algorithms, 2017, CoRR [abs/1710.08665](https://arxiv.org/abs/1710.08665). arXiv:1710.08665. URL: <http://arxiv.org/abs/1710.08665>.
- [32] LLC Gurobi Optimization, Gurobi optimizer reference manual, 2020, URL: <http://www.gurobi.com>.
- [33] Matthew Roughan, Simplifying the synthesis of internet traffic matrices, SIGCOMM Comput. Commun. Rev. (ISSN: 0146-4833) 35 (5) (2005) 93–96, <http://dx.doi.org/10.1145/1096536.1096551>.
- [34] Srikanth Sundaresan, Cristian Lumezanu, Nick Feamster, Pierre Francois, Autonomous traffic engineering with self-configuring topologies, in: Proceedings of the ACM SIGCOMM 2010 Conference, in: SIGCOMM '10, Association for Computing Machinery, New Delhi, India, ISBN: 9781450302012, 2010, pp. 417–418, <http://dx.doi.org/10.1145/1851182.1851239>.
- [35] L.G. Valiant, A scheme for fast parallel communication, SIAM J. Comput. 11 (2) (1982) 350–361, <http://dx.doi.org/10.1137/0211027>, eprint: <https://doi.org/10.1137/0211027>.