

On a New Competitive Measure for Oblivious Routing

Gábor Németh

Dept. of Telecommunications and Media Informatics
Budapest University of Technology and Economics
Budapest, Hungary

Abstract—Oblivious routing algorithms use only locally available information at network nodes to forward traffic, and as such, a plausible choice for distributed implementations. It is a natural desire to quantify the performance penalty we pay for this distributedness. Recently, Räcke has shown that for general undirected graphs the *competitive ratio* is only $O(\log n)$, that is, the maximum congestion caused by the oblivious algorithm is within a logarithmic factor of the best possible congestion. And while the performance penalty is larger for directed networks (Azar gives a $\Omega(\sqrt{n})$ lower bound), experiments on many real-world topologies show that it usually remains under 2. These competitive measures, however, are of worst-case type, and therefore do not always give adequate characterization.

The more different combinations of demands a routing algorithm can accommodate in the network without congestion, the better. Driven by this observation, in this paper we introduce a new competitive measure, the *volumetric competitive ratio*, as the measure of all admissible demands compared to the measure of demands routed without congestion. The main result of the paper is a general lower bound on the volumetric ratio; and we also show a directed graph with $O(1)$ competitive ratio that exhibits $\Omega(n)$ volumetric ratio.

Our numerical evaluations show that the competitiveness of oblivious routing in terms of the new measure quickly vanishes even in relatively small common-place topologies.

Keywords—*competitive ratio; oblivious routing; ℓ^p norm; L^p norm; throughput polytope; feasible region; probability of congestion; hyper-spherical coordinates*

I. INTRODUCTION

Routing algorithms are used to drive the process of forwarding traffic from source nodes to remote destination nodes through a network. Often, network links are of limited capacity and it is also the task of the routing algorithm to ensure that no network link gets seriously overloaded with excess traffic. Minimizing congestion, however, may require the global knowledge of the actual traffic demand pattern the users pose to the network, which is very difficult to ensure in a distributed setting. The class of non-adaptive routing algorithms with the property that only knowledge that is locally available at network nodes is used when making forwarding decisions is called *oblivious routing algorithms*. In fact, in oblivious routing a set of paths and corresponding splitting ratios are precomputed offline, which are then applied in the on-line phase statically to the incoming traffic at network nodes. This scheme is easy to

implement in a distributed fashion, both in virtual-circuit-based as well as in a packet routed environment.

The performance of an oblivious routing algorithm is usually described in terms of the maximum congestion it produces, compared to the congestion produced by an optimal routing algorithm [1]. The maximization is taken over all possible demand combinations. Surprisingly, this *competitive ratio* (or oblivious performance ratio [2]) in undirected networks is only logarithmic [1], [3], [4], and even though no such appealing characterization exists for directed networks, experiment suggests that it rarely surpasses 2 [5]. This makes oblivious routing an attractive choice for implementing distributed routing algorithms.

Being a worst-case measure, the competitive ratio, however, might not always give adequate statistical representation of the performance penalty of distributed routing, as compared to optimal routing. In this paper, we study the competitive ratio arising as the extent to which an oblivious routing algorithm can route without congestion demands that otherwise could be routed in the network by a properly chosen, possibly centralized optimal routing algorithm.

This work is connected to the scientific program of the "Development of quality-oriented and cooperative R+D+I strategy and functional model at BME" project. This project is supported by the New Hungary Development Plan (Project ID: TÁMOP-4.2.1/B-09/1/KMR-2010-0002).

Consider the following formal definition. Given a capacitated graph G and an arbitrary routing algorithm, let the feasible region D be the set of demands that can be routed by the algorithm without congestion in G . Let the throughput polytope T of G be the set of demands that can be routed without congestion at all (i.e., the feasible region of an optimal algorithm). Then, the quantity of *all* demands routable in the network is the volume $V(T)$, the quantity of demands routed by the routing algorithm without congestion is $V(D)$, and their ratio, the so called *volumetric competitive ratio* $\alpha_V = \frac{V(T)}{V(D)}$ represents the fraction of routable demands the algorithm can handle successfully. Easily, $\alpha_V > 1$, and the smaller α_V the better. Sometimes, we use the measure $\alpha_{POC} = 1 - \frac{V(D)}{V(T)}$ as it has appealing statistical interpretation: provided that demands arrive from the set of routable demands T according to a uniform distribution, what is the chance that we encounter congestion at some link in the network. In other words, α_{POC} quanti-

fies the frequency that an operator can expect to find his network in a congested state, in the case when absolutely no information on traffic demands is available *a priori*. Thus, $\alpha_{POC} = 1 - \frac{1}{\alpha_v}$ is called the *probability of congestion (POC)*.

In this paper, we study the competitiveness of oblivious routing in terms of the new performance measures.

A. Related Work

Oblivious routing on hypercubes was first studied by Valiant and Brebner [6]. They give a randomized oblivious routing scheme with $O(\log n)$ competitive ratio. The first oblivious routing algorithm for generic undirected graphs is due to Räcke [1], with the remarkable property that the maximum congestion is within a $O(\log^3 n)$ factor of the lowest possible congestion, attainable by an optimal algorithm (for such an optimal adaptive algorithm, see [7]). This result was non-constructive, as the algorithm's off-line running time was exponential. The competitive ratio was subsequently improved by Harrelson, Hildrum and Rao in [3] to $O(\log^2 n \log \log n)$. They also show a polynomial algorithm for constructing the hierarchical tree decomposition that underlies Räcke's oblivious routing scheme (see also [8]). Finally, it was also Räcke who was able to improve the competitive ratio to $O(\log n)$ in [4]. This bound is asymptotically tight as there are very simple networks (e.g., two dimensional grids) on which no oblivious routing algorithm exists with sub-logarithmic competitive ratio [9], [10].

For directed graphs, unfortunately, no logarithmic congestion guarantee exists. Azar et al. give a $\Omega(\sqrt{n})$ lower bound in [2]. To cope with this intrinsic difficulty, Hajiaghayi et al. [11] present oblivious routing schemes that achieve $O(\log^2 n)$ competitive ratio provided that demands arrive randomly from a known demand distribution. Even though one might expect this assumption to improve the upper bound in undirected networks too, this is not the case: Hajiaghayi et al. give a $\Omega\left(\frac{\log n}{\log \log n}\right)$ lower bound in [12].

The first polynomial algorithm to obtain the best oblivious routing scheme for specific input graphs was introduced by Azar in [2]. A simpler, linear programming-based algorithm was given in [5]. Here, the task is to, given a directed or undirected graph as input, compute the static routing that produces the smallest competitive ratio possible on this graph. This can (and usually is) better than logarithmic. For instance, extensive numerical evaluations suggest that the competitive ratio in most real-world network topologies remains under 2 [5]. Unfortunately, this approach is not suitable to obtain the generic upper bound Räcke could obtain with the use of hierarchical tree decompositions (whose approach, in turn, does not yield optimal oblivious routing schemes for particular graphs).

Demand for more descriptive performance measures for oblivious routing has increased lately [13] [14]. The motivation is not necessarily to quantify the performance of oblivious routing algorithms, but rather to drive the optimization algorithms that compute them. For instance, one would better consider the average network load, or the sum-of-squares of the loads, as the performance measure, in contrast to the maximum load.

Following on the work of Gupta [15], Engler and Räcke in [16] give a universal treatment, able to treat the above case and many more. They define a generic aggregation function that determines how loads at individual links are converted to a congestion measure for the network, and then show a $O(\log n)$ -competitive oblivious routing algorithm when the aggregation function is an ℓ^p norm. Their development is non-constructive, which was recently remedied by Bhaskara and Vijayaraghavan [17]. Note, however, that these performance measures are still of worst-case nature, meaning that it is the maximum of the congestion measure experienced over all possible demands that determines the outcome.

It seems that Rétfvári et al. were the first ones to systematically study the geometric properties of the feasible region D and the throughput polytope T [18], [19]. They showed that, under reasonable regularity conditions, both D and T are compact, down-monotone, K -dimensional polyherda (K is the number of source-destination pairs). They also showed that no polynomial-size description exists for T even in very small networks. Thus, T is usually given implicitly, in the form of a linear program.

For computing the volumetric ratio, we need to obtain the volume of D and T . Unfortunately, except very low dimensions or special polytopes with high degree of symmetry (e.g., simplices, hyper-cubes), this is a very hard task. In particular, Elekes showed that one cannot construct a general polynomial-time algorithm for calculating the volume of K -dimensional bodies [20]. Therefore, randomized algorithms were proposed to break down the complexity and approximate the volume with a prescribed absolute/relative error [21]-[23]. These algorithms rely on Monte-Carlo integration and introduce random walks for sampling. The complexity of the best known randomized method is $O(K^4)$ linear program solver calls [24]. Even though linear programs can be solved in polynomial time, this still can be prohibitive in large networks. What is worse, random-walk-based sampling is another significant source of complexity [25], as a linear program needs to be solved at each step of the random walk, and we need thousands of random samples obtained in possibly thousands of steps.

B. Our Results

In this paper, we introduce a new competitive measure of non-worst-case type to better characterize the performance of distributed oblivious routing as compared to optimal centralized routing. The measure, called the volumetric competitive ratio, quantifies the fraction of routable demands an algorithm can handle without congestion.

In the first part, we give performance bounds on oblivious routing in terms of the new competitive measure. For this, we develop a geometric model, using which for directed graphs we give a $\Omega(n)$ worst-case lower bound on the volumetric ratio. This behavior is exhibited even in cases when the standard competitive ratio is $O(1)$. Then, we obtain a universal upper bound for the volumetric ratio. At the moment, it is not known whether these bounds are tight.

In the second part of the paper, we conduct brief numerical evaluations on real-world topologies, which indicate that the

measures quickly exhibit the worst case behavior as K increases.

II. NOTATIONS AND DEFINITION

Let $G = (V, E)$ be a connected directed or undirected graph ($n = |V|$ and $m = |E|$), with positive edge capacities $c_e: e \in E$, and K source-destination pairs $(s_k, d_k) \in V \times V$. A traffic matrix (or, simply, demand) is a column K -vector $\theta = [\theta_k: k = 1, \dots, K]$, where θ_k represents the request of the k -th source-destination pair. Let S be a routing algorithm which, given some demand θ , generates a flow $g_k(e)$ (i.e., a routing) for each $k = 1, \dots, K$ on each edge $e \in E$. For our purposes, it is enough to know that the output of S is an aggregate flow on each edge e : $S(\theta) = [f(e) = \sum_k g_k(e): e \in E]$. We denote the relative flow (or load) $f(e)/c(e)$ produced by S on e by $L_S(e)$.

Gupta et al. introduce the notion of aggregation functions $\text{agg}: \mathbb{R}^E \rightarrow \mathbb{R}$ to aggregate the loads of individual edges into a cost measure [15]. Engler and Racke in [16] study the case when the aggregation function is an ℓ^p norm $\|L_S\|_p = (\sum_{e \in E} |L_S(e)|^p)^{\frac{1}{p}}$. Then, they give oblivious routing algorithms to minimize the *competitive ratio* α_p defined by this ℓ^p norm:

$$\alpha_p = \max_{\mathbb{R}_+^K} \frac{\|L_{\text{OBL}}\|_p}{\|L_{\text{OPT}^p}\|_p} = \left\| \frac{\|L_{\text{OBL}}\|_p}{\|L_{\text{OPT}^p}\|_p} \right\|_{L^\infty(\mathbb{R}_+^K)}. \quad (1)$$

Here, OPT^p is an optimal routing algorithm that for each demand θ assigns the flow L_{OPT^p} that minimizes the ℓ^p norm $\|L_{\text{OPT}^p}\|_p$, and L_{OBL} is the flow produced for θ by the oblivious routing algorithm that minimizes (1). The second term in (1) comes from substituting max by the L^∞ norm (for a good overview on ℓ^p and L^p norms, see [26]). Easily, different settings of p yield different interesting algorithms. For $p = \infty$ specifically we get the well-known ‘‘congestion-minimizing’’ oblivious routing algorithms [1]-[4]. In this paper, we mean by ‘‘oblivious routing’’ this very case (i.e., when $p = \infty$), but throughout the developments we shall often use different settings for p . Similarly, the term ‘‘competitive ratio’’ will mean (1) with choosing $p = \infty$, i.e., α_∞ .

Our task in this paper is to seek alternatives to the above competitive measure. Our approach is mainly geometric, the main ingredients of which are as follows. Given a routing algorithm S , let the *feasible region* D of S be the set of demands θ to which S orders a routing that does not violate edge capacities:

$$D = \{\theta: f(e) \leq c(e) \ \forall e \in E, \text{ where } [f(e)] = S(\theta)\}.$$

The throughput polytope T of G is the set of all routable demands, i.e., the feasible region of OPT^∞ . Under the above assumptions, both D and T are K -dimensional, compact, convex, down-monotone polytopes [18]. Therefore both sets are measurable in terms of the standard Lebesgue measure, that is, the K -dimensional volumes $V(D)$ and $V(T)$ exist and are non-zero. Note that we call a set X *down-monotone*, if $x \in X \Rightarrow \forall 0 \leq y \leq x: y \in X$.

We shall need some definitions from geometry to deal with polytope volumes. Let P denote a convex, compact down-monotone polytope in \mathbb{R}_+^K with $0 \in P$. Let ϕ_P denote the gauge functional of P , i.e., $\phi_P(x) = \min\{\eta \in \mathbb{R}: x \in \eta P\}$ for all $x \in \mathbb{R}_+^K$. Note that ϕ_P is a spherical function, that is, $\forall \alpha \in \mathbb{R}: \phi_P(\alpha x) = \alpha \phi_P(x)$. The reason why we invoke the gauge functional is that it is a natural geometric generalization of the competitive ratio $\alpha_\infty: \alpha_\infty = \max_{\theta \in T} \phi_D(\theta)$. Using this notation, the volume of P is given by [23], [27]

$$V(P) = \frac{1}{K!} \int_{\mathbb{R}_+^K} e^{-\phi_P(x)} d\mu(x), \quad (2)$$

where μ is the measure function on \mathbb{R}^K . Converting the integration to hyper-spherical coordinates one can rewrite the above from integration on the entire positive orthant to integration on the surface of the unit K -ball ∂B :

$$V(P) = \frac{1}{K} \int_{\partial B} \phi^{-K}(x) d\mu(x) = \frac{1}{K} \int_{\partial B} d_p^K(x) d\mu(x), \quad (3)$$

where $d_p(x) = \max_{x \in \partial B} \{\eta \in \mathbb{R}: \eta x \in P\}$ denotes the distance of the boundary point of P from the origin in the direction defined by the point $x \in \partial B$. For notational convenience, we shall often omit the dependence on x and simply write $d\mu$ for $d\mu(x)$ and d_p for $d_p(x)$.

III. THE VOLUMETRIC COMPETITIVE RATIO

As mentioned previously, in this paper we want to characterize the fraction of routable demands a routing algorithm can handle without congestion. Consider the following definition:

Definition 1: Given a capacitated network G with K source-destination pairs and a routing algorithm S , let T be the throughput polytope of G and let D be the feasible region of S in G . Then, the *volumetric competitive ratio* of algorithm S is defined as

$$\alpha_V = \frac{V(T)}{V(D)}, \quad (4)$$

and the Probability of Congestion (POC) is $\alpha_{\text{POC}} = 1 - \frac{V(D)}{V(T)} = 1 - \frac{1}{\alpha_V}$.


The reason of why we also define the POC is that it has relevant practical interpretation: α_{POC} quantifies the chance that we find the network in a congested state, when traffic demands arrive from T according to a uniform distribution.

In the next sections, we search global bounds on the above competitive measures. First, we discuss directed graphs and then we turn to generic bounds on undirected networks.

A. A Worst-case Upper Bound in Directed Graph

First, we show that there exist directed graphs that exhibit $\Omega(n)$ volumetric competitive ratio, even though the standard competitive ratio α_∞ is $O(1)$.

Theorem 1: For any $n \geq 4$, there is a directed graph of n nodes with $\alpha_\infty < 2$ and $\alpha_V > \frac{n}{6}$.



Proof: Consider the directed graph in Fig. 1 for any $K \geq 2$,
Fig. 1. Directed graph (a) for illustrating the proof of Theorem 1, and the corresponding throughput polytope T and the feasible region D of the oblivious routing for $K = 2$ (b). For this case, $\alpha_\infty = 4/3$ and $\alpha_V = 3/4$.

let all link capacities be 1, and let the source-destination pairs be $(s_k, K + 2): s_k \in \{1, \dots, K\}$. First, we construct the oblivious routing function w.r.t. to the ℓ^p norm and calculate $\alpha_\infty(K)$ as a function of K , then we obtain an approximation on α_V .

1. *Competitive ratio:* Construct the oblivious routing function as follows. The first source-destination pair has only a single path, thus all its traffic is sent through this path. The rest of the users $k \in \{2, \dots, K\}$ have two paths. Due to the symmetry of the network, it is enough to consider a general user k . Let β denote the fraction of traffic sent at s_k to the path passing through node $K + 1$. We consider two critical classes of traffic matrices, as these will produce the largest link load for oblivious routing. One is $[1, 1, \dots, 1]^T \in T$, for which the maximum load $1 + \sum_{k=2}^K \beta$ occurs on link $(K + 1, K + 2)$. Second, for some general $k \in \{2, \dots, K\}$ the demand $[0, 0, \dots, 2, \dots, 0]^T \in T$, where there is a single non-zero element in the k -th position, causes $2(1 - \beta)$ maximum load on link $(k, K + 2)$. To find $\alpha_\infty(K)$, we seek for β so that the maximum load is minimal. This occurs exactly when $1 + \sum_{k=2}^K \beta = 2(1 - \beta)$, which yields $\alpha_\infty(K) = \frac{2K}{K+1} \leq 2$, for every finite K .

2. *The volumetric ratio:* We need the volume of T and D . Instead of calculating the volumes directly, we take lower and upper approximations. Let T' and D' be such that $T' \subseteq T$ and $D' \supseteq D$. By down-monotonicity of T , the K -hypercube of size 1 resides completely inside T .

Moreover, we find $K - 1$ half-hypercubes in T as well, where the k -th half-hypercube, $k \in \{2, \dots, K\}$, is obtained by placing a K -hypercube of size 1 at the point $[0, \dots, 1, \dots, 0]^T$ (1 is in the k -th position) and taking the intersection with the half-

space $\theta_1 + \theta_k \leq 2$. The volume of T' is the sum of the volumes of the above polytopes: $V(T') = 1 + \frac{1}{2}(K - 1)$. Second, we give an outer approximation D' for D . Simply put, D is enclosed by a hyper-rectangle, whose lower left corner is the origin and whose upper right corner is the point $[1, 1 + \frac{1}{K}, \dots, 1 + \frac{1}{K}]^T$. Let this hyper-rectangle be D' and so $V(D') = (1 + \frac{1}{K})^{K-1}$.

Putting all together, we get the desired $\alpha_V = \frac{V(T')}{V(D')} \geq \frac{V(T')}{V(D')} = \frac{1 + \frac{1}{2}(K-1)}{(1 + \frac{1}{K})^{K-1}} \geq \frac{n}{6}$ with the substitution $n = K + 2$. ■

B. A Worst-case Lower Bound for Undirected Graphs

The result of the previous section is only valid for a class of special directed networks, in which obtaining good approximations on $V(D)$ and $V(T)$ is possible. For general graphs, this approach is not viable. Therefore, we shall pursue a different approach below: we give a general lower bound on α_V in terms of α_∞ . Consider the following theorem.

Theorem 2: $\alpha_V \geq \frac{V(T)}{\sqrt{V(T')}} (\alpha_\infty)^{\frac{K}{2}}$,

where $T' = \{y \in \mathbb{R}_+^K: \exists x \in T, y = x^2\}$.

First, consider the following technical Lemma.

Lemma 1: Let $f \leq g$ and $g^{2K}, (\frac{f}{g})^{2K}$ belong to $L^2(\partial B_r)$ and let $\mu(\partial B_r) = 1$. Then

$$\frac{\int_{\partial B_r} f^K d\mu}{\int_{\partial B_r} g^K d\mu} \leq \frac{\sqrt{\int_{\partial B_r} g^{2K} d\mu}}{\int_{\partial B_r} g^K d\mu} \sqrt{\int_{\partial B_r} \left[\frac{f}{g}\right]^{2K} d\mu}.$$

Lemma 1 is the direct consequence of the Cauchy—Schwarz—Bunyakovsky inequality [26], [28], [29].

Proof of Theorem 1: Using (3) and the above Lemma, we get the following upper bound on $\frac{1}{\alpha_V}$.

$$\begin{aligned} \frac{1}{\alpha_V} &= \frac{V(D)}{V(T)} = \frac{\int_{\partial B} d_D^K d\mu}{\int_{\partial B} d_T^K d\mu} = \frac{\int_{\partial B_r} d_D^K d\mu}{\int_{\partial B_r} d_T^K d\mu} \\ &\leq \frac{\sqrt{\int_{\partial B_r} d_T^{2K} d\mu}}{\int_{\partial B_r} d_T^K d\mu} \sqrt{\int_{\partial B_r} \left[\frac{d_D}{d_T}\right]^{2K} d\mu} \\ &= \frac{\sqrt{V(T')}}{V(T)} \sqrt{\left\| \frac{\|L_{OPT^\infty}\|_\infty}{\|L_{OBL}\|_\infty} \right\|_{L^{2K}(\partial B_r)}^K} \end{aligned} \tag{5}$$

where the radius $r: \mu(\partial B_r) = 1$. The last equation comes from observing that the term $\frac{d_D}{d_T}$ is the value the feasible region should be scaled to enclose all the points of the throughput polytope in the selected direction. Thus,

$$\frac{d_T}{d_D} = \frac{\|L_{OBL}\|_\infty}{\|L_{OPT^\infty}\|_\infty}. \tag{6}$$

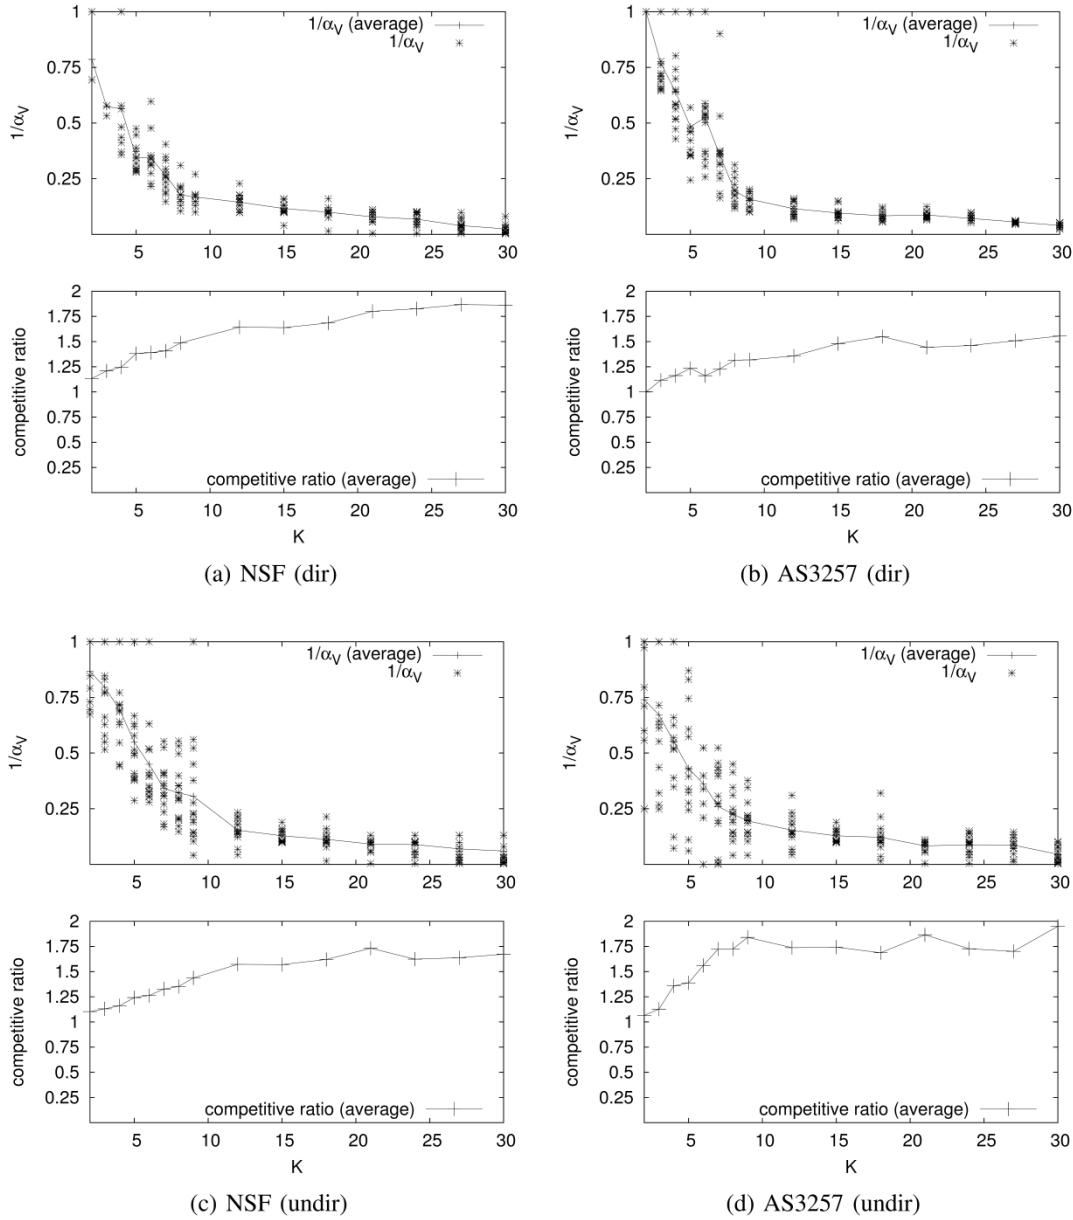


Fig. 2. The competitive ratio α_∞ and the inverse of the approximated volumetric ratio α_V w.r.t. the number of source-destination pairs K for selected directed and undirected networks.

We rewrite (5) so that the outer norm is a L^∞ norm. First, we use fact that $\lim_{q \rightarrow \infty} \|f\|_{L^q(\partial B_r)} = \|f\|_{L^\infty(\partial B_r)}$. Second, choosing let r so that let $\mu(\partial B_r) = 1$ yields $\forall q \in \mathbb{N}: \|f\|_{L^{q+1}(\partial B_r)} \geq \|f\|_{L^q(\partial B_r)}$. Thus, we write (5) as:

$$\alpha_V = \frac{\int_{\partial B_r} a_D^K d\mu}{\int_{\partial B_r} a_T^K d\mu} \geq \frac{V(T)}{\sqrt{V(T')}} \left(\frac{\|L_{OPT^\infty}\|_\infty^K}{\|L_{OBL}\|_\infty \|L^{2K}(\partial B_r)\|} \right)^{-1/2} \geq \frac{V(T)}{\sqrt{V(T')}} \left(\lim_{q \rightarrow \infty} \frac{\|L_{OPT^\infty}\|_\infty^K}{\|L_{OBL}\|_\infty \|L^{2K+q}(\partial B_r)\|} \right)^{-1/2}$$

Using Jensen's inequality [26] for the concave function $\frac{1}{id_{\mathbb{R}}}$ yields:

$$\frac{\int_{\partial B_r} a_D^K d\mu}{\int_{\partial B_r} a_T^K d\mu} \geq \frac{V(T)}{\sqrt{V(T')}} \sqrt{\lim_{q \rightarrow \infty} \frac{1}{\frac{\|L_{OPT^\infty}\|_\infty^K}{\|L_{OBL}\|_\infty \|L^{2K+q}(\partial B_r)\|}}} \geq \frac{V(T)}{\sqrt{V(T')}} \sqrt{\lim_{q \rightarrow \infty} \frac{\|L_{OBL}\|_\infty^K}{\|L_{OPT^\infty}\|_\infty \|L^{2K+q}(\partial B_r)\|}}$$

Therefore

$$\frac{\int_{\partial B_r} d_D^K d\mu}{\int_{\partial B_r} d_T^K d\mu} \geq \frac{V(T)}{\sqrt{V(T')}} \sqrt{\lim_{q \rightarrow \infty} \left\| \frac{\|L_{OBL}\|_\infty}{\|L_{OPT}\|_\infty} \right\|_{L^\infty(\partial B_r)}^K},$$

which completes the proof.

IV. NUMERICAL EVALUATIONS

Finally, we seek the answer for the question whether the worst case bounds on the volumetric ratio presented in Section III indeed appear in real networks. Therefore, we conducted some numerical evaluations using the volume approximation algorithm of the previous section. We ran the evaluations on the ISP data maps from the Rocketfuel dataset [30]. We used the same method as in [5] to obtain approximate POP-level topologies: we collapsed the topologies so that nodes correspond to cities, we eliminated leaf-nodes and we set link capacities inversely proportional to the link weights. In this paper, only results for the network AS3257 are shown, as we observed similar results for the rest of the networks as well. Another round of evaluations was conducted on the NSFNET Phase II topology [31]. From these topologies, we generated two series of increasingly more complex networks by adding gradually more source-destination pairs¹. Recall that the number of source-destination pairs K determines the dimension of the underlying geometric space, and so it has profound impact on α_V . In our experiments, K was increased from 2 to 30. For each K , fifteen independent samples were generated picking the source and the destination nodes randomly according to a bimodal distribution and then α_∞ and α_V were evaluated for each scenario. The parameters were chosen so that the result has larger than 10% relative error with less than 5% probability. Fig. 2 depicts the results for both networks. Note that we show $\frac{1}{\alpha_V}$ instead of α_V for better visualization.

Our most important observation is that real networks indeed exhibit the worst-case behavior seen in Section III. We observe that with the increase of K , $\frac{1}{\alpha_V}$ rapidly approaches 0. This suggests that already in networks with more than a couple of source-destination pairs it is only a very small fraction of all the routable demands the oblivious routing algorithm can handle without congestion. In terms of the probability of congestion, which approaches 1 as $\frac{1}{\alpha_V}$ approaches 0, this basically means that a network adopting oblivious routing will spend most of its time in a congested state, provided that demands arrive uniformly from the set an optimal algorithm could route without any congestion at all. And this is despite of the fact that α_∞ remains low (we see $\alpha_\infty < 2$ in both networks for all K).

V. CONCLUSION

Oblivious routing is a promising candidate for minimum-congestion routing in large networks. This is thanks to that, on the one hand, it is a fundamentally distributed scheme and, on the other hand, it comes equipped with a hard performance guarantee, namely, the maximum congestion it causes is within a logarithmic factor of the best possible congestion. This performance characterization, however, is intrinsically of worst-case nature.

In this paper, we introduced an alternative competitive measure, the so called volumetric ratio, which measures the fraction of routable demands an oblivious routing algorithm can route without congestion. We observed that already in very small directed networks (i.e., the ones in Section III-A), oblivious routing algorithms order infeasible routing to $O\left(\frac{1}{n}\right)$ fraction of the, otherwise routable, demands. We showed further worst-case bounds valid for both directed and undirected graphs.

A disadvantage of the new measure is that it is very difficult to numerically evaluate it. In fact, negative results on exact polytope volume computation suggest that we cannot hope for a polynomial time algorithm to compute α_V .

Finally, by numerical evaluations we showed that the worst-case behavior we identified clearly manifests itself in real networks.

Easily, this paper is only a first step towards a more thorough performance characterization of oblivious routing. At the moment, it is unclear whether our bounds are tight, and we seriously lack proper upper bounds. It seems though that, with some more work, the geometric model we introduced in this paper will be able to provide these results.

REFERENCES

- [1] H. Räcke, "Minimizing congestion in general networks," in Proceedings of the 43rd Symposium on Foundations of Computer Science, FOCS '02, pp. 43–52, 2002.
- [2] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke, "Optimal oblivious routing in polynomial time," in Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, STOC '03, (New York, NY, USA), pp. 383–388, ACM, 2003.
- [3] C. Harrelson, K. Hildrum, and S. Rao, "A polynomial-time tree decomposition to minimize congestion," in Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures, SPAA '03, (New York, NY, USA), pp. 34–43, ACM, 2003.
- [4] H. Räcke, "Optimal hierarchical decompositions for congestion minimization in networks," in Proceedings of the 40th annual ACM symposium on Theory of computing, STOC '08, pp. 255–264, 2008.
- [5] D. Applegate and E. Cohen, "Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs," in ACM SIGCOMM, pp. 313–324, 2003.
- [6] L. G. Valiant and G. J. Brebner, "Universal schemes for parallel communication," in Proceedings of the thirteenth annual ACM symposium on Theory of computing, STOC '81, pp. 263–277, 1981.
- [7] G. Rétvári and G. Németh, "On optimal multipath rate-adaptive routing," in IEEE Symposium on Computers and Communications, ISCC'10, pp. 605–610, IEEE Computer Society, 2010.
- [8] M. Bienkowski, M. Korzeniowski, and H. Räcke, "A practical algorithm for constructing oblivious routing schemes," in Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures, SPAA '03, (New York, NY, USA), pp. 24–33, ACM, 2003.
- [9] Y. Bartal and S. Leonardi, "On-line routing in all-optical networks," in Proceedings of the 24th International Colloquium on Automata, Languages and Programming, ICALP '97, pp. 516–526, 1997.
- [10] B. Maggs, F. Meyer auf der Heide, B. Vocking, and M. Westermann, "Exploiting locality for data management in systems of limited bandwidth," in Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on, pp. 284–293, oct 1997.
- [11] M. Hajiaghayi, J. Kim, T. Leighton, and H. Räcke, "Oblivious routing in directed graphs with random demands," in STOC '05, pp. 193–201, 2005.
- [12] M. T. Hajiaghayi, R. D. Kleinberg, T. Leighton, and H. Räcke, "New lower bounds for oblivious routing in undirected graphs," in Proceed-

A preliminary version of the results were presented as a poster at ACM Sigmetrics 2012 [32].

- ings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, SODA '06, (New York, NY, USA), pp. 918–927, ACM, 2006.
- [13] P. Harsha, T. Hayes, H. Narayanan, H. Räcke, and J. Radhakrishnan, “Minimizing average latency in oblivious routing,” in SODA '08, (Philadelphia, PA, USA), pp. 200–207, 2008.
- [14] G. Lawler and H. Narayanan, “Mixing times and L_p bounds for oblivious routing,” in Workshop on Analytic Algorithmics and Combinatorics, (ANALCO'09) 4, 2009.
- [15] A. G. M. T. Hajiaghayi and H. Räcke, “Oblivious network design,” in In Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 970–979, 2006.
- [16] M. Englert and H. Räcke, “Oblivious Routing for the L_p -norm,” IEEE Foundations of Computer Science, pp. 32–40, 2009.
- [17] A. Bhaskara and A. Vijayaraghavan, “Computing the matrix p -norm,” CoRR, vol. abs/1001.2613, 2010.
- [18] G. Rétvári, J. J. Bíró, and T. Cinkler, “Fairness in capacitated networks: A polyhedral approach,” in IEEE INFOCOM, vol. 1, pp. 1604–1612, May 2007. G. Rétvári and G. Németh, “Demand-oblivious routing: distributed vs. centralized approaches,” in Proceedings of the 29th conference on Information communications, INFOCOM'10, (Piscataway, NJ, USA), pp. 1217–1225, IEEE Press, 2010.
- [19] G. Elekes, “A geometric inequality and the complexity of computing volume,” Discrete and Computational Geometry, vol. 1, pp. 289–292, 1986.
- [20] M. Dyer, A. Frieze, and R. Kannan, “A random polynomial-time algorithm for approximating the volume of convex bodies,” J. ACM, vol. 38, pp. 1–17, January 1991.
- [21] M. Dyer and A. Frieze, “Computing the Volume of Convex Bodies: A Case where Randomness Provably Helps,” in Proceedings of Symposia in Applied Mathematics, vol. 44, 1991.
- [22] L. Lovász and M. Simonovits, “Random walks in a convex body and an improved volume algorithm,” Random Structures & Algorithms, vol. 4, no. 4, pp. 359–412, 1993.
- [23] L. Lovász and S. Vempala, “Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm,” J. Comput. Syst. Sci., vol. 72, pp. 392–417, March 2006.
- [24] R. Kannan, L. Lovász, and M. Simonovits, “Random walks and an $O^*(n^5)$ volume algorithm for convex bodies,” Random Structures & Algorithms, vol. 11, no. 1, pp. 1–50, 1997.
- [25] G. B. Folland, Real Analysis: Modern Techniques and Their Applications. Pure and Applied Mathematics, John Wiley & Sons, Inc., 2 ed., 1999.
- [26] P. Gritzmann and V. Klee, “On the Complexity of Some Basic Problems in Computational Convexity: II. Volume and mixed volumes,” in Polytopes: Abstract, Convex and Computational (T. Bisztriczky, P. McMullen, R. Schneider, and A. W. Weiss, eds.), vol. 440 of NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., pp. 373–466, Dordrecht: Kluwer Acad. Publ., 1994.
- [27] H. Royden and P. Fitzpatrick, Real Analysis. Prentice Hall, 4 ed., 2007.
- [28] W. Rudin, Real and Complex Analysis. International Series in Pure and Applied Mathematics, McGraw-Hill, 3 ed., 1986.
- [29] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, “Inferring link weights using end-to-end measurements,” in ACM IMC, pp. 231–236, 2002.
- [30] B. Chinoy and H. W. Braun, “The National Science Foundation network.” Tech. Rep., CAIDA, available online: <http://www.caida.org/outreach/papers/1992/nsfn/nsfn-t1-technology.pdf>, Sep 1992.
- [31] G. Németh and G. Rétvári, “Towards a statistical characterization of the competitiveness of oblivious routing (poster).” ACM Sigmetrics, 2012.