Towards a Statistical Characterization of the Competitiveness of Oblivious Routing

Gábor Németh and Gábor Rétvári

{nemethgab, retvari}@tmit.bme.hu

HSNLab, Dept. of Telecommunications and Media Informatics, BME

Motivation

Oblivious routing asks for a static routing that serves arbitrary user demands with minimal performance penalty. Performance is measured in terms of the competitive ratio, the proportion of the maximum congestion to the best possible congestion. In undirected networks, the competitive ratio is upper-bounded by a logarithmic function of the number of nodes, and its value usually remains under 2 in directed networks. **Is the oblivious routing really that good? Or the competitive measure hides some crucial details on how oblivious routing performs?**

Model

The **throughput polytope** T is the set of traffic matrices θ for which there is a routing u that accommodates θ in the network with no link over-utilization:

$$T = \{\theta : \exists u \geq 0 \text{ so that } \sum_{P \in \mathcal{P}_k} u_P = \theta_k \quad \forall k \in \mathcal{K}$$
$$\sum_{k \in \mathcal{K}} P_k u_k \leq c\}$$

For a static routing function $\mathcal{S}$, the feasible region $R(\mathcal{S})$ is a down-monotone a polytope:

$$R(\mathcal{S}) = \{\theta : \sum_{k \in \mathcal{K}} P_k f_k \theta_k \leq c,\ \theta \geq 0\}\ . \qquad (1)$$

New Performance Metrics

The **probability of congestion** (PoC) η is the probability that at least one network link gets overloaded by oblivious routing:

$$\eta = 1 - \frac{\mathrm{Vol}(R)}{\mathrm{Vol}(T)}\ , \qquad (2)$$

where Vol denotes the K-dimensional volume.

Given a static routing function $\mathcal{S}$, the **expected value of congestion** (EVoC) is the mean value of the maximum link utilization produced by $\mathcal{S}$:

$$\mu(\mathcal{S}) = \mathbb{E}\left[\kappa_{\mathcal{S}}(\theta)\right]\ . \qquad (3)$$

Approximating EVoC

The EVoC can be approximated as

$$\frac{K}{K+1}\left(\frac{1}{1-\eta}\right)^{1/K} \leq \mu \leq \alpha\ , \qquad (4)$$

where K is the number of source-destination pairs.

References

[1] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke: *Optimal oblivious routing in polynomial time*, STOC 2003

[2] H. Räcke: *Minimizing congestion in general networks*, FOCS 2002

[3] G. Rétári and G. Németh: *Demand-oblivious routing: distributed vs. centralized approaches*, IEEE INFOCOM 2010

Example

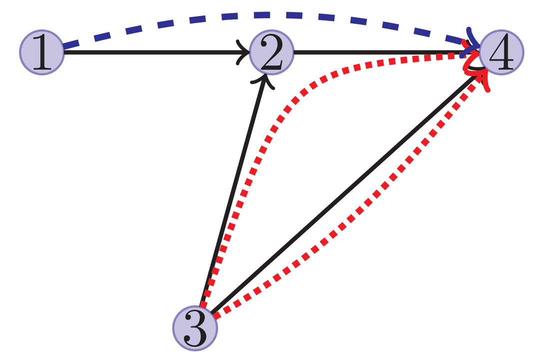

$(s_1, d_1) = (3, 4)$

$(s_2, d_2) = (1, 4)$

$P_1 = \{(3,4)\}$

$P_2 = \{(3,2),(2,4)\}$

$P_3 = \{(1,2),(2,4)\}$

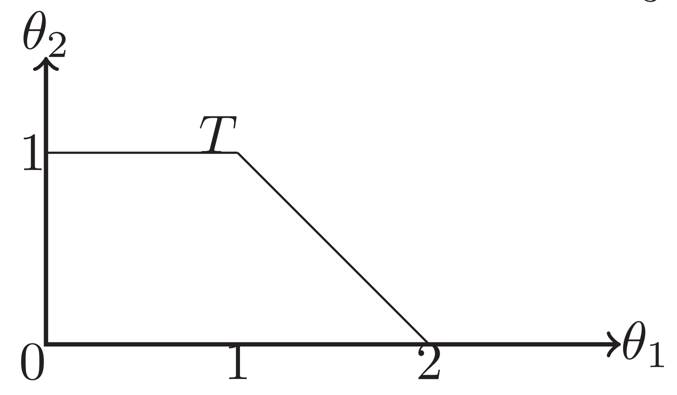

$$\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = f_1\theta_1 = \begin{pmatrix} \frac{2}{3} \\ \frac{1}{3} \end{pmatrix}\theta_1, \qquad u_3 = f_2\theta_2 = \theta_2\ .$$

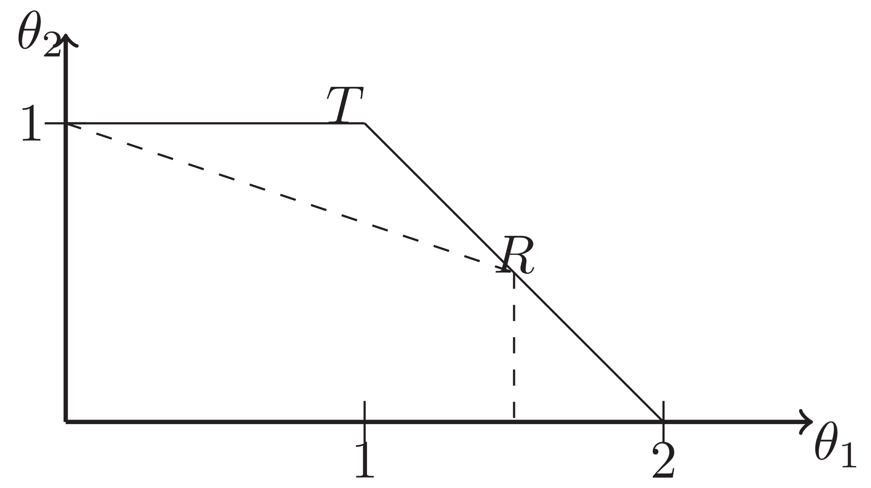

PoC and EVoC

- our most important observation is that the probability of congestion grows beyond 80% for as few as about 10 users, and as the number of users enters the range of the number of nodes, and it approaches 100% with very high confidence
- the expected value of congestion grows beyond 1, indicating symptoms of grave congestion

NSFNET Phase II (directed and undirected)

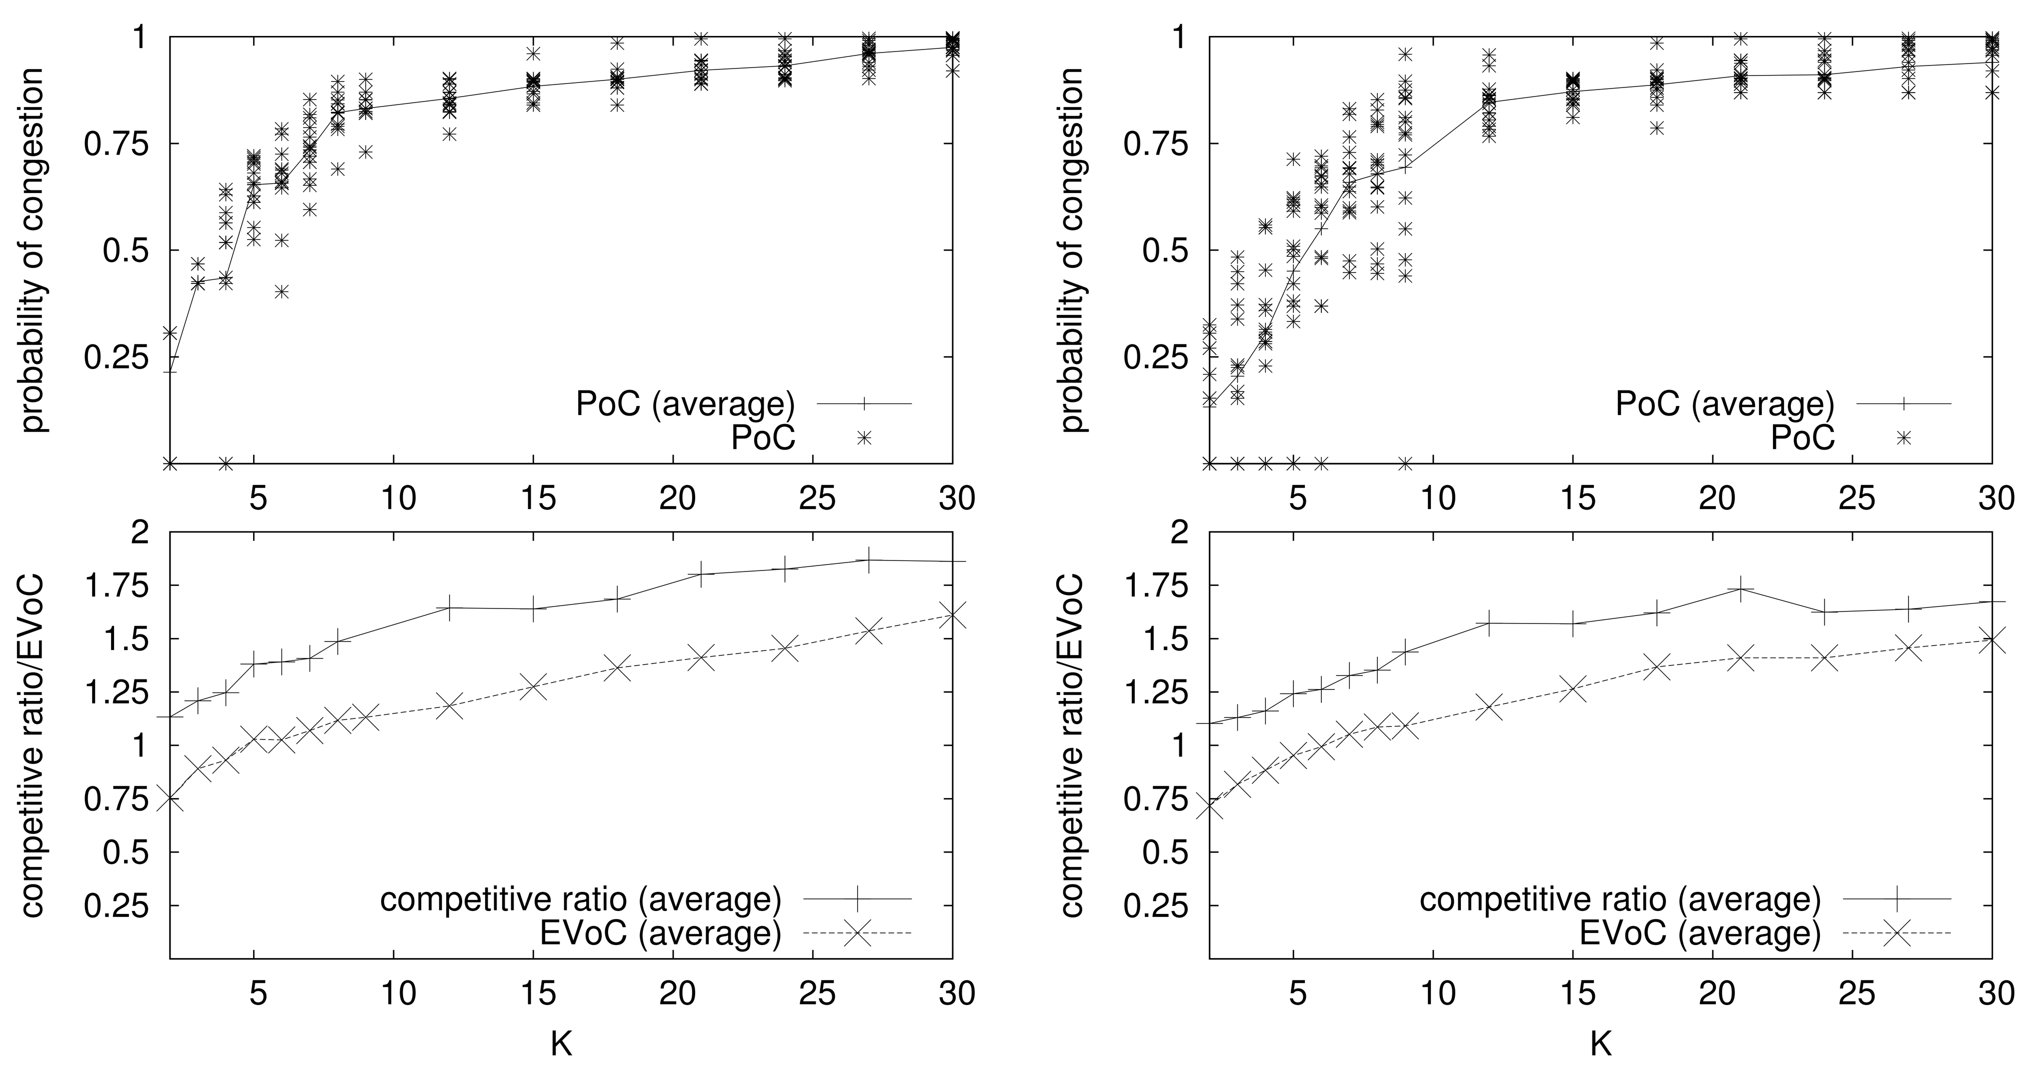

AS 3257 directed and undirected

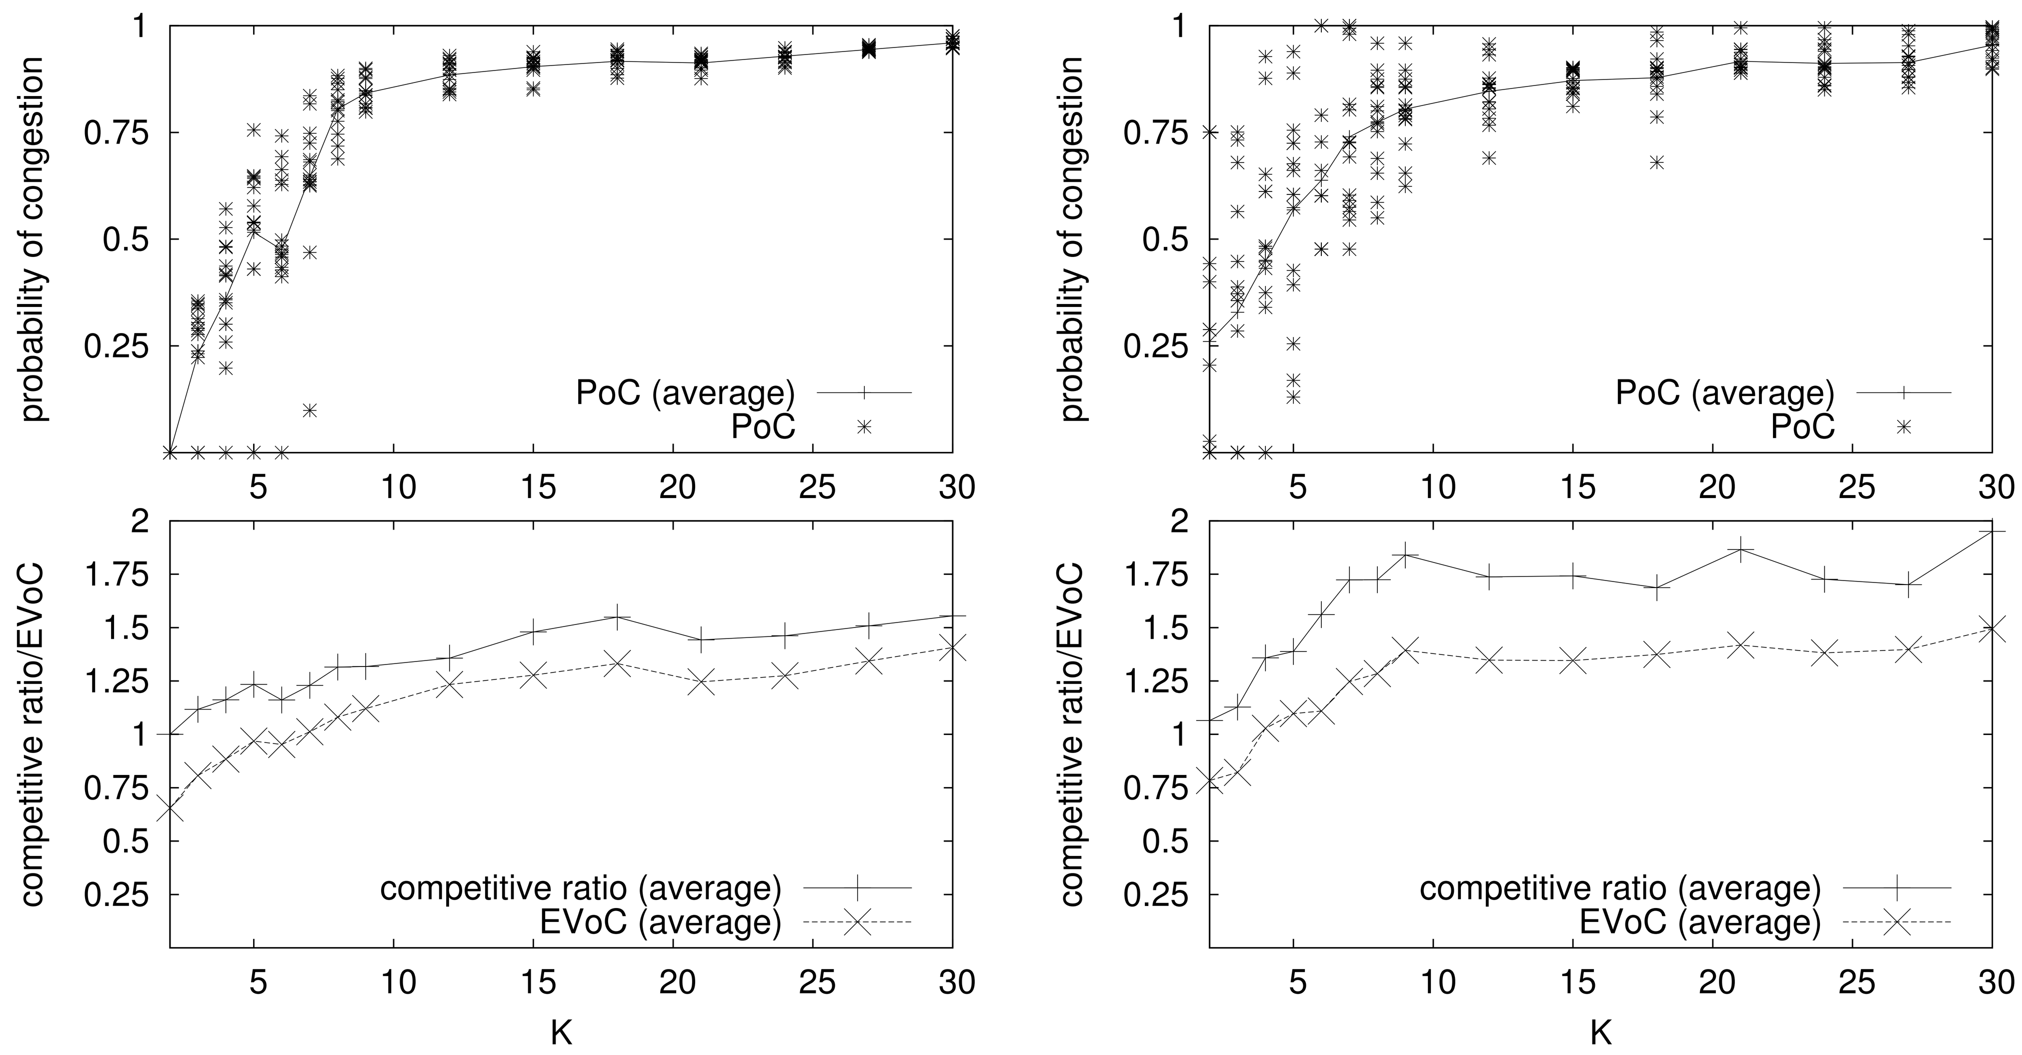

PoC: $K \leq 9$ exact results were computed using Vinci; $K > 9$ approx. result (the chance of relative error larger than 10% is less than 10%)

Future Works

- more exhaustive numerical evaluations
- theoretical lower bounds
- fast approx. algorithm for PoC