Self-organizing weights for Internet AS-graphs and surprisingly simple routing metrics

Jan Scholz and Martin Greiner

Self-organizing weights for Internet AS-graphs and surprisingly simple routing metrics

Jan Scholz[1,2] and Martin Greiner[1]

[1]Aarhus School of Engineering and Department of Mathematical Sciences,
Aarhus University, Ny Munkegade 118, DK-8000 Aarhus C, Denmark
[2]Frankfurt Institute for Advanced Studies and Frankfurt International Graduate School
for Science, Johann Wolfgang Goethe Universität, Ruth-Moufang-Straße 1,
60438 Frankfurt am Main, Germany

Abstract

The transport capacity of Internet-like communication networks and hence their efficiency may be improved by a factor of 5–10 through the use of highly optimized routing metrics, as demonstrated previously. Numerical determination of such routing metrics can be computationally demanding to an extent, that prohibits both investigation of and application to very large networks. In an attempt to find a numerically less expensive way of constructing a metric with a comparable performance increase, we propose a local, self-organizing iteration scheme and find two surprisingly simple and efficient metrics. The new metrics have negligible computational cost and result in an approximately 5-fold performance increase, providing distinguished competitiveness with the computationally costly counterparts. They are applicable to very large networks and easy to implement in today's Internet routing protocol on the AS-level.

Introduction

Today's society relies heavily on communication networks, exemplified by their most prominent specimen, the Internet. It provides the substrate for a wealth of activities, ranging from email communication to on-demand video broadcasting, to name but a few. The ubiquity of networked systems and their far reaching implications for everyday life in our modern world have sparked a broad interest of science (Watts and Strogatz, 1998; Albert et al., 1999; Newman, 2001; Barabási, 2009). The quest for optimized utilization of the network connections to enhance communication flow within networks is one of obvious significance (Tadić et al., 2007) and has particularly attracted the attention of the physics community (Guimerà et al., 2002; Braunstein et al., 2003; Zhao et al., 2005; Danila et al., 2006; Yan et al., 2006; Scholz et al., 2008; Yang et al., 2009). Like many other networks found in various fields of science,

---

Corresponding author: Jan Scholz, scholz@imf.au.dk

1

from biology (Yook et al., 2004) to social sciences (Granovetter, 2003) and communication networks in general, the Internet exhibits a complex structure (Faloutsos et al., 1999), far from the homogeneous configuration that is described by classical random graphs (Erdös and Rényi, 1960). A common feature of these networks is the small-world characteristic (Milgram, 1967; Watts and Strogatz, 1998), in many cases accompanied by a scale-free degree distribution (Faloutsos et al., 1999; Albert et al., 1999; Newman, 2001; Ebel et al., 2002). In the past decade, tremendous effort has been put forth to explain the emergence of these complex networked systems and to understand their dynamics (Barabási, 2009).

1   Traffic model and routing metrics

To estimate the performance of communication networks consider a network defined by the graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with $N = |\mathcal{V}|$ nodes and $M = |\mathcal{E}|$ links. In the considered traffic model data packets are created at a rate μ at every node, each with a random destination. As today's Internet routers operate in full duplex mode we assume bandwidth limits on the links, but unlimited processing capacity at the nodes. This corresponds to $\texttt{M/M/}k$ queuing in the standard queuing theory nomenclature (Kendall, 1953). For simplicity and because these quantities are notoriously difficult to determine for the real Internet (Krioukov et al., 2007), we assume uniform bandwidths on the links and, without loss of generality, set these to 1 for both directions, i.e. each link e_{ij} has the capacity to carry 1 packet per time step from node i to j and another one from j to i. Each directed link e has to carry a certain amount of traffic load L_e. The link $e_{\mathrm{BN}} = \arg\max_e L_e$ which carries the highest load is the bottleneck of the network and limits its transport capacity T_{e2e}. In analogy to the case of limited node bandwidth and unlimited link bandwidth (Guimerà et al., 2002; Zhao et al., 2005), we get the total end-to-end transport capacity

$$T_{\mathrm{e2e}} = \mu_{\mathrm{crit}} N = \frac{N(N-1)}{L_{e_{\mathrm{BN}}}} \tag{1.1}$$

of the network (Scholz et al., 2008). The critical rate μ_{crit} is the packet generation rate at which the bottleneck link is utilized exactly at its bandwidth and all other links are utilized below their bandwidth capacity.

In communication networks employing deterministic routing, data packets follow the shortest route with respect to a routing metric w, that assigns a weight w_e to the link $e \in \mathcal{E}(\mathcal{G})$. These weights determine the length of a path $\mathcal{P}$, given by the sum of the weights on its links:

$$\mathrm{length}(\mathcal{P}, w) := \sum_{e \in \mathcal{E}(\mathcal{P})} w_e \,. \tag{1.2}$$

From the set $\mathrm{paths}(i, f)$ of all paths from source i to target f, the shortest paths according to the metric w are

$$\mathrm{spath}(i, f; w) := \underset{\mathcal{P} \in \mathrm{paths}(i,f)}{\arg\min} \ \mathrm{length}(\mathcal{P}, w) \,. \tag{1.3}$$

Should there be multiple shortest paths, one of these paths is chosen with uniform probability to route a packet.

For the case of shortest path routing according to metric w the load of link e is given by:

$$L_e := \sum_{i,j \in \mathcal{V}} \frac{|\{\mathcal{P} \in \mathrm{spath}(i,j;w) : e \in \mathcal{P}\}|}{|\mathrm{spath}(i,j;w)|} \,. \tag{1.4}$$

The terms in the sum correspond to the probability, for a packet with source i and destination j, to be routed over e. For the case of unweighted links, $w_e = 1$ for all $e \in \mathcal{E}(\mathcal{G})$, the resulting load L_e is known as the so called *betweenness centrality* (Freeman, 1977; Brandes, 2001).

Compared to unweighted shortest path routing, which is the de facto standard in current routing on the Internet's AS level, previous work (Fukś and Lawniczak, 1999; Yan et al., 2006; Krause et al., 2006; Schäfer et al., 2006; Danila et al., 2006; Scholz et al., 2008) has found optimized metrics that are able to increase the transport capacity T_{e2e} by a factor of 5–10, depending on the underlying network structure. The gist of the most successful metrics, the *extremal* metric (Danila et al., 2006) and the *hybrid* metric (Scholz et al., 2008), is to find the highest loaded link of the network, the bottleneck e_{BN}, and to increase its weight, $w_{e_{\mathrm{BN}}}$, either additively or multiplicatively. As the optimization applies only to the maximally loaded link, it is a form of extremal optimization (Danila et al., 2006). Because the weights influence the shortest paths, which in turn determine the load and hence the bottleneck, the procedure has to be iterated. As successful as the application of extremal optimization is, it is also a weak point of the metrics. In every iteration step, only a single weight is adjusted, which necessitates a relatively large number of steps, compared to for example the *smoothing* metric (Schäfer et al., 2006), which updates all link weights in each iteration. The number of iteration steps makes the extremal and the hybrid metric numerically hard to obtain, requiring numerical calculations that scale with the number of nodes N approximately like $\mathcal{O}\left(N^2 \log N\right)$ (Scholz et al., 2008) or worse $\mathcal{O}\left(N^3 \log N\right)$ (Danila et al., 2006). Here we use the hybrid metric as a benchmark for new, computationally less demanding metrics. In an attempt to find numerically less expensive metrics with similar performance increases, we investigate a local, self-organizing weight assignment and find a simple and surprisingly efficient weight assignment.

2 Self-organizing weights

Looking for a way to reduce the number of needed iterations, we introduce a rule, that every node can apply by itself, in concert with its neighbors, using information as local as possible, e.g. it must not be necessary to determine a global maximum of the load. Additionally, we demand that the weights are conserved by the local rule, i.e. weights shall not be created or destroyed, but only moved from one link to another. For a stepwise update of the link weights w we propose the following rule,

which we refer to as the *self-organizing (SO) metric*:

$$w_m^{t+1} = \frac{1}{2} \sum_{n \in \mathcal{L}(m)} \left(\frac{w_m^t}{\ell_m} + \frac{w_n^t}{\ell_n} \right)$$

$$+ \frac{\epsilon}{2} \sum_{n \in \mathcal{L}(m)} \left(\frac{w_m^t}{\ell_m} + \frac{w_n^t}{\ell_n} \right) \text{sign}(L_m^t - L_n^t), \qquad (2.1)$$

where $\mathcal{L}(m)$ is the link neighborhood of link m, ℓ_m is its link degree, and $\text{sign}(x)$ is the signum function.[1] The weight of link m at the iteration time t is given by w_m^t, and L_m^t denotes the load of link m subject to the weights at time t. The effect of the first sum in Equation (2.1) is to average the weights in the link neighborhood, while the second sum results in a shift of weights from links with a relatively low load to neighbors that are relatively high loaded. The influence of the second sum in relation to the first is controlled by the parameter ϵ, which may take values $0 \leq \epsilon < 1$. Weight conservation is fulfilled by the SO metric, as can easily be verified by calculating $\sum_{m \in \mathcal{E}} w_m^{t+1}$. The contributions of the first sum in Equation (2.1) sums up to $\sum_m w_m^t$, and the second sum vanishes, as for every instance of the signum function, another one with the negative argument exists. From this $\sum_{m \in \mathcal{E}} w_m^{t+1} = \sum_{m \in \mathcal{E}} w_m^t$ follows. As a change of the weights results in different packet routing, the link loads L_e have to be recalculated after the weight update.

Starting from initial weights set to 1 for every link, Figure 1 shows the development of the end to end transport capacity T_{e2e} of an AS-level Internet snapshot (CAIDA Macroscopic Topology Project Team, 2000–2006) during the iteration of the update rule. After a sufficient number of iteration steps, the network's performance becomes stationary around a mean value. The transport capacity in this stationary state depends on the parameter ϵ, with the optimal performance of the weight update (2.1) achieved by $\epsilon \approx 0.35$. That said, this stationary state optimum by no means guarantees an improvement of T_{e2e} relative to the hop metric, and in fact, for the concrete network realization shown in Figure 1 it does not. However, Figure 1 highlights a peculiarity for smaller ϵ-values, especially for $\epsilon = 0$. Here T_{e2e} shows a pronounced peak around the 2nd to 3rd step. Focusing on this peculiar case, we further investigate the iteration of Equation (2.1) for $\epsilon = 0$.

---

[1]The signum function is defined as $\text{sign}(x) = -1$ for $x < 0$, $\text{sign}(x) = 0$ for $x = 0$, and $\text{sign}(x) = 1$ for $x > 0$.
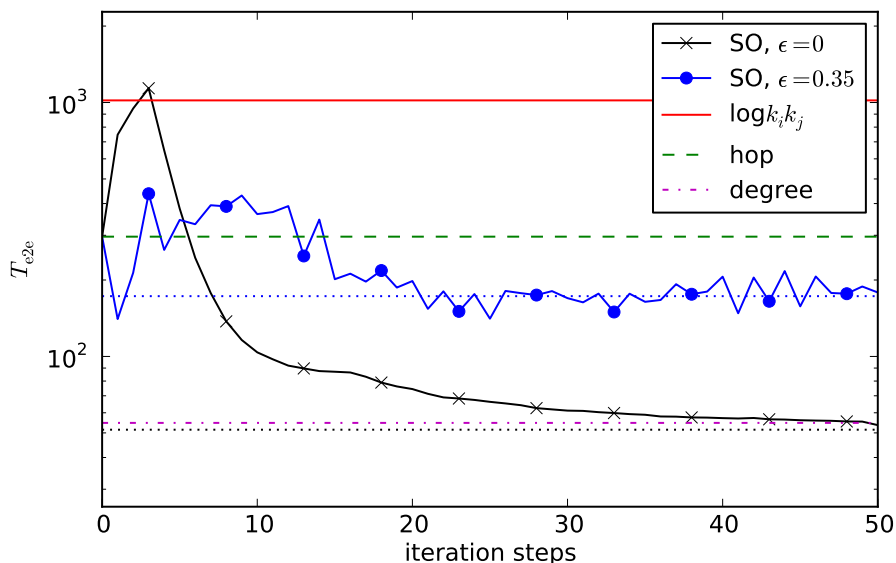
Figure 1: Evolution of the transport capacity T_{e2e} during iteration of the SO metric (2.1) with $\epsilon = 0$ and $\epsilon = 0.35$. The dotted horizontal lines indicate the asymptotic T_{e2e} in the steady state of the iteration, determined by averaging the values from step 80 to 100. Dashed horizontal lines show T_{e2e} achieved by the three static weight assignments hop metric, node degree metric (Yan et al., 2006), and, in anticipation of our results, the $\log(k_i k_j)$ metric. The network topology used here is an AS-scan from June 2006, with 8111 nodes and 22370 links (CAIDA Macroscopic Topology Project Team, 2000–2006).

3 The iteration's asymptotic state

For the case of $\epsilon = 0$ we can determine the resulting link weights in the limit $t \to \infty$ by interpreting Equation (2.1) as a difference equation, and introducing Δt, which had been implicitly set to 1 before. We get

$$\frac{w_m^{t+1} - w_m^t}{\Delta t} = -\frac{1}{2} w_m^t + \frac{1}{2} \sum_{n \in \mathcal{L}(m)} \frac{w_n^t}{\ell_n}. \tag{3.1}$$

Using matrix notation and going to infinitesimal time steps, this can be written as

$$\frac{\partial w(t)}{\partial t} = -\frac{1}{2} \left(\boldsymbol{I} - AD^{-1} \right) w(t), \tag{3.2}$$

where $\boldsymbol{I}$ is the identity matrix, A is the adjacency matrix of the line graph $\widehat{\mathcal{G}}(\mathcal{G})$ of $\mathcal{G}$, and D is the diagonal matrix of the link degrees ℓ_m. Abbreviating $\boldsymbol{I} - AD^{-1} = \tilde{\Lambda}$ the differential equation is solved by

$$w(t) = \exp\left(-\tfrac{1}{2} \tilde{\Lambda} t \right) w(0). \tag{3.3}$$

Using the eigendecomposition $\tilde{\Lambda} = \tilde{V} \lambda \tilde{V}^{-1}$, this can be written as:

$$w(t) = \tilde{V} \exp\left(-\tfrac{1}{2} \lambda t \right) \tilde{V}^{-1} w(0), \tag{3.4}$$

5

where λ is the diagonal matrix of eigenvalues λ_i of $\tilde{\Lambda}$, and $\tilde{V}$ is the matrix of corresponding eigenvectors $\tilde{v}_i$. As, through the similarity transform $D^{\frac{1}{2}}$, the linear operator $\tilde{\Lambda}$ is similar to the normalized graph Laplacian (Chung, 1997)

$$\Lambda = \boldsymbol{I} - D^{-\frac{1}{2}}AD^{-\frac{1}{2}}, \tag{3.5}$$

the eigenvalues of $\tilde{\Lambda}$ and Λ are identical and for connected graphs the eigenvalues λ_i of $\tilde{\Lambda}$ are all positive, except $\lambda_0 = 0$. Hence, in the limit $t \to \infty$ all eigenvectors are damped out by the matrix exponential in Equation (3.4), except $\tilde{v}_0$ corresponding to λ_0. Using the similarity transform $D^{\frac{1}{2}}$, $\tilde{v}_0$ can be calculated from v_0, the zeroth eigenvector of the normalized Laplacian:

$$\tilde{v}_0 = D^{\frac{1}{2}}v_0. \tag{3.6}$$

As the normalized Laplacian's zeroth eigenvector is given by $v_0 = D^{\frac{1}{2}}\mathbf{1}$, see (Chung, 1997), for $\epsilon = 0$ the link weight

$$\lim_{t\to\infty} w_m^t \sim \ell_m \tag{3.7}$$

becomes proportional to the link degree independent of the initial weight configurations.
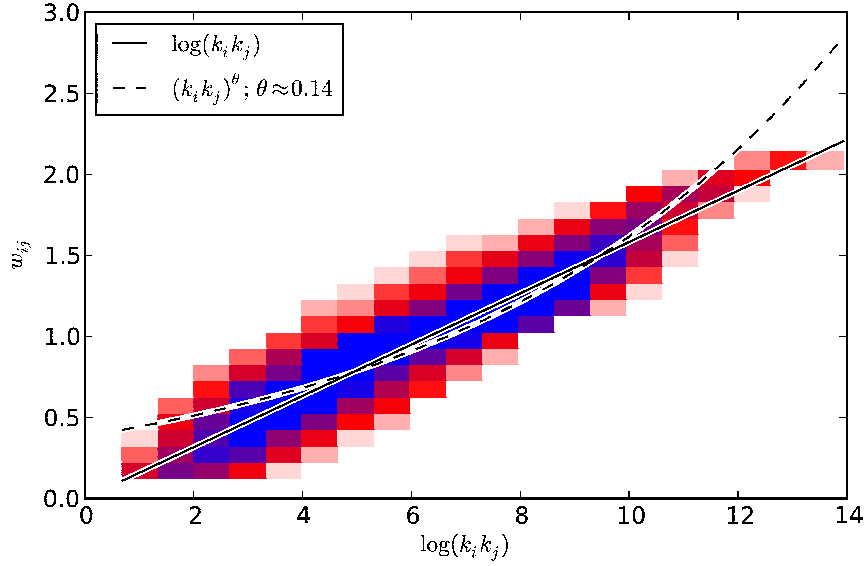
4 The iteration's peak

Using the weight distribution $p(w_{ij}|k_i k_j)$ we analyze the weight assignment that leads to the peak in the performance measured by T_{e2e} around the 2nd to 3rd iteration step of Equation (2.1) with $\epsilon = 0$. The weight distribution accounts for the probability of a link e_{ij} to be assigned weight w_{ij}, given the product $k_i k_j$ of the adjacent nodes' degrees. The weight distribution directly at the peak is shown in Figure 2a. It exhibits a well defined shape, that shows a clear relation of the weight and the product of node degrees

$$w_{ij}(k_i, k_j) \sim \log(k_i k_j), \tag{4.1}$$

shown by a least squares fit in Figure 2a. Another recently discussed (Yang et al., 2009) functional form for link weights is

$$w_{ij}(k_i, k_j) \sim (k_i k_j)^\theta, \tag{4.2}$$

which approximates the weight distribution at the SO-metric's peak for $\theta \approx 0.14$ almost as good as $\log(k_i k_j)$. The sum of squared deviations of fitting $N_{\mathrm{l}} \log(k_i k_j)$ to the weight distribution amounts to 80% of fitting $N_{\mathrm{p}}(k_i k_j)^\theta$. For comparison, the asymptotic state of the iteration with $\epsilon = 0.35$ is shown in Figure 2b. Unlike the case for $\epsilon = 0$, it exhibits no apparent structure and is therefore not easily approximated by a simple function of the degrees.

(a) Weight distribution at the $\epsilon = 0$ peak.
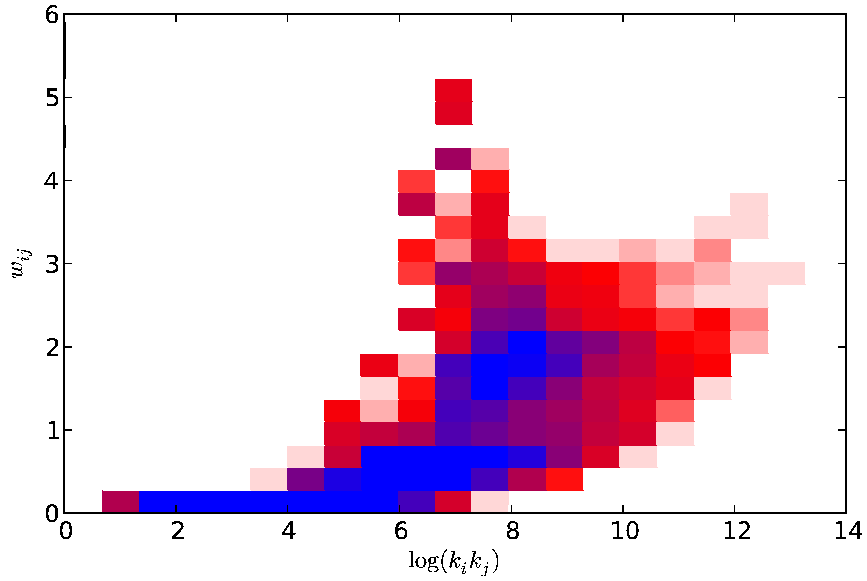
(b) Asymptotic weight distribution for $\epsilon = 0.35$

Figure 2: Weight distributions produced by the SO metric for (a) $\epsilon = 0$ at the performance peak during iteration, and (b) for the asymptotic state with $\epsilon = 0.35$. The topology is the same as in Figure 1.

The surprisingly simple relations between link weight and node degrees found in Figure 2a urge to pragmatically give the weight assignments

$$w_{ij} = \log(k_i k_j) \; , \tag{4.3}$$

which we refer to as the $\log(k_i k_j)$-metric, and the $(k_i k_j)^{0.14}$-metric

$$w_{ij} = (k_i k_j)^{\theta} \; ; \; \theta = 0.14 \tag{4.4}$$

a try.

Figure 3 shows the transport capacity T_{e2e} of these metrics compared to the hop metric, the smoothing metric (Schäfer et al., 2006) adapted to link weights (Scholz et al., 2008), the $(k_i k_j)^{0.4}$-metric as proposed to enhance the resilience to cascading failures of nodes (Yang et al., 2009), and most notably the hybrid metric (Scholz et al., 2008). The hybrid metric continues to be the best of the tested metrics, but
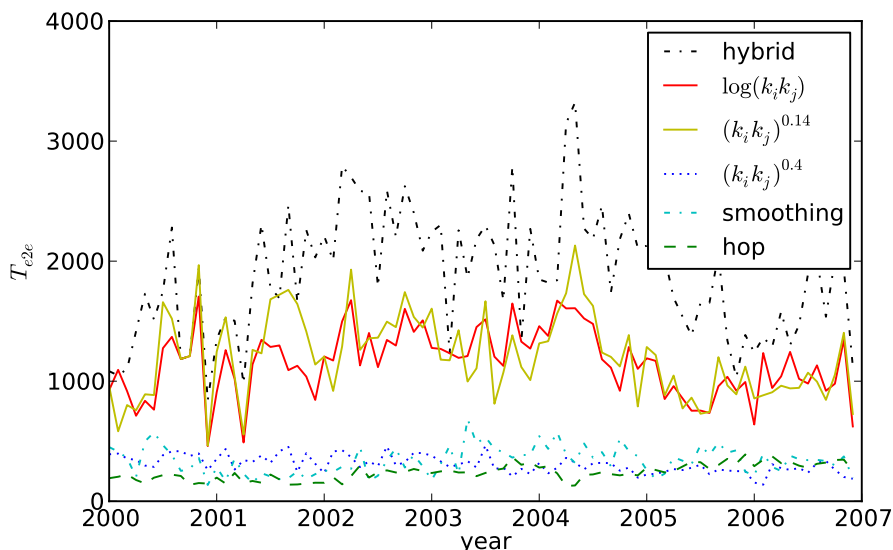

Figure 3: Comparison of the transport capacity T_{e2e} achieved by the $\log(k_i k_j)$ metric and other metrics on monthly Internet snapshots by the CAIDA project. The $\log(k_i k_j)$-metric and the $(k_i k_j)^{0.14}$-metric are compared to the hop, smoothing (Scholz et al., 2008), $(k_i k_j)^{0.4}$ (Yang et al., 2009), and hybrid metric (Scholz et al., 2008). The node degree metric and the link degree metric (asymptotic state of the SO metric with $\epsilon = 0$, see Equation (3.7)) generally perform similar to the hop metric or worse, but are left out, to avoid cluttering the figure.

the performances of the $\log(k_i k_j)$-metric and the $(k_i k_j)^{0.14}$-metric are of the same order of magnitude. Relative to the hop metric both result in gain ratios $g \approx 5$ which compare very well with $g \approx 7$ for the hybrid metric. The $(k_i k_j)^{0.4}$-metric, determined as optimal with respect to cascading node failures and transport capacity in the case of limited node capabilities (Yang et al., 2009), can not guarantee an increase of the transport capacity with respect to limited link capabilities investigated here.

The huge performance increase achieved by the two metrics determined from Figure 2a is especially impressive when comparing the computational complexity to

the computational complexity of the hybrid metric. As given by Scholz et al. (2008), the computational complexity of iterating the hybrid metric is $\mathcal{O}\left(N^2 \log N\right)$. In contrast the $\log(k_i k_j)$-metric and $(k_i k_j)^{0.14}$-metric do not need to be iterated, instead they merely determine the degrees of nodes and calculate the logarithm for every node, hence the complexity is $\mathcal{O}\left(N\right)$. Assuming the calculation is done by every node itself, it is trivial to parallelize, in which case the complexity reduces to $\mathcal{O}\left(1\right)$.

# 5 Sensitivity to topology changes

To investigate whether the success of the $\log(k_i k_j)$ metric is confined to the special network structure of the AS-level Internet we gradually perturb the structure of the given network. The perturbations are done in two ways, to which we will refer as *randomization* and *rewiring*.

As *randomization* we denote the crosswise exchange of links, where two random links e_{ab} and e_{cd} between distinct nodes a, b, c, d are removed and the links e_{ad} and e_{cb} are added instead, assuring that these links have not been present before. This modification of the network structure conserves the degree of all involved nodes and therefore the overall degree distribution as well. We quantify the progression of randomization, by the ratio of exchanged links M_{xchg} to total number of links M in the network. For sufficient progression of randomization $M_{\text{xchg}} \gg M$, the resulting network is equivalent to a realization of the configuration model (Newman, 2003) with the degree distribution of the unperturbed network. The effect of randomization on the T_{e2e} performance is shown in Figure 4a. The node degree metric and the $\log(k_i k_j)$ metric both profit from the shortcuts gradually added by the randomization process. The transport capacity of the hop metric on the other hand is essentially unaffected. Throughout the whole progression of randomization, the $\log(k_i k_j)$ metric maintains an increase of transport capacity of a factor of 5 or more compared to the hop metric.

We refer to the second structure perturbation scheme as *rewiring*. Here a random link is deleted, followed by the addition of a link between two random, not yet connected nodes. This variant of structure perturbation has for example been employed in (Watts and Strogatz, 1998) to make the transition from regular grids to random graphs. Like in the case of randomization, we quantify the progression of rewiring by the fraction of rewired links to the total number of links in the network. For sufficiently large progression of rewiring $M_{\text{rwr}} \gg M$, the network structure approaches a Poissonean graph, with an average degree $\langle k \rangle$ and size N given by the original graph. The effect of rewiring on T_{e2e} is shown in Figure 4b. Again the transport capacity achieved by the degree and $\log(k_i k_j)$ metric is continuously increased by the introduced shortcuts. The hop metric shows a slight decrease of T_{e2e} for a rewiring progression up to approximately $M_{\text{rwr}} \gtrsim 2M$, with a steep ascent for larger rewiring progressions. Figure 5 offers an explanation for this effect, as it shows how the scale-free property of the last remaining high degree nodes, so called hubs, is removed by the rewiring procedure around $M_{\text{rwr}} \approx 2M$. The node degree and $\log(k_i k_j)$ metric do not suffer from overloaded remaining hubs, as both metrics explicitly avoid nodes with high degree. Once the rewired network structure is in the regime of random graphs, the effect of the different metrics vanishes, consistent
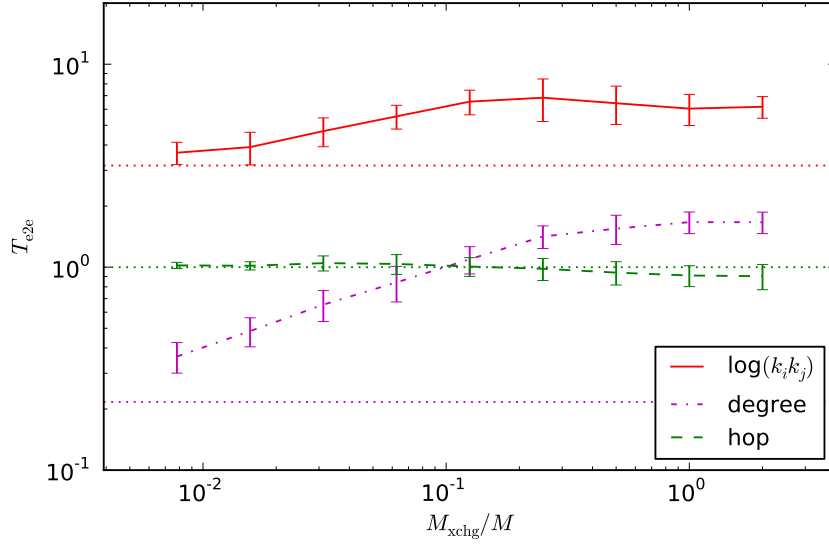
(a) Randomization
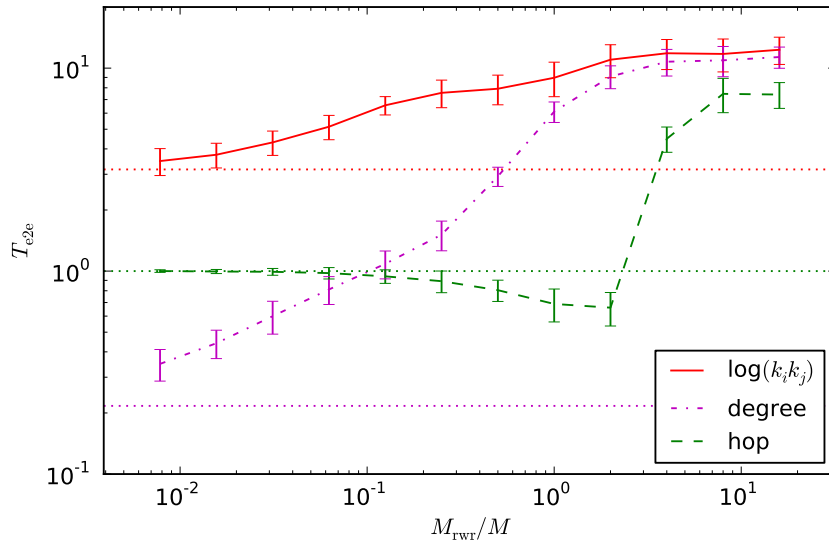
(b) Rewiring

Figure 4: Change of T_{e2e} based on the hop, node degree, and $\log(k_i k_j)$ metric under (a) randomization and (b) rewiring.

10

with prior results (Scholz et al., 2008).

The connection of optimal transport efficiency to resilience to cascading failures, as implicitly used by previous studies (Schäfer et al., 2006; Scholz et al., 2008) and explicitly stated by Yang et al. (2009), suggests to employ the $\log(k_i k_j)$-metric in the context of protection against cascading failures (Motter and Lai, 2002) in future studies. Here the numerical simplicity of the $\log(k_i k_j)$-metric is of extraordinary importance, as the change of network topology during the cascade steps demands a timely re-determination of the weights.
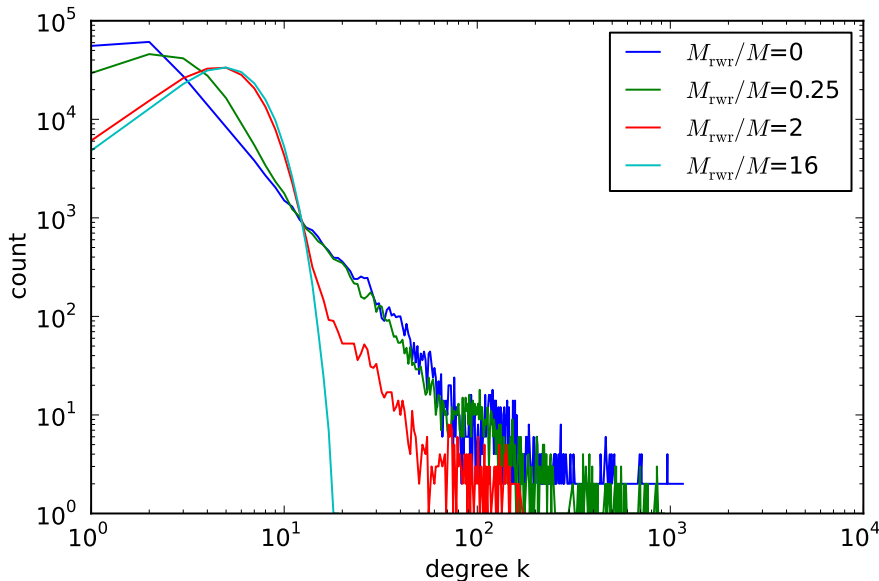

Figure 5: Change of degree distribution $p(k)$ under rewiring.

6   Conclusion

The investigation of a self-organized, local link weight assignment rule has lead to the discovery of the simple, but surprisingly efficient $\log(k_i k_j)$-metric and $(k_i k_j)^{0.14}$-metric. The enhancement of transport capacity achieved through application of the $\log(k_i k_j)$ metric proves to be robust with respect to the topology changes in the sense, that in general it results in an approximately 5-fold performance increase relative to the hop metric. Compared to previously proposed metrics with comparable efficiency, the advantage of both newly found metrics is the simplicity with respect to numerical complexity. While other efficient metrics demand considerable computational resources, the $\log(k_i k_j)$-metric and $(k_i k_j)^{0.14}$-metric are computable in no time. It is thus straightforward to implement these metrics, for example in the Internets routing protocol on the AS-level.

Acknowledgments

References

Réka Albert, Hawoong Jeong, and Albert-László Barabási. Internet: Diameter of the world-wide web. *Nature*, 401(6749):130–131, September 1999. ISSN 0028-0836. doi: 10.1038/43601.

Albert-László Barabási. Scale-free networks: A decade and beyond. *Science*, 325(5939):412–413, 2009. doi: 10.1126/science.1173299.

Ulrik Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001. doi: 10.1080/0022250X.2001.9990249.

Lidia A. Braunstein, Sergey V. Buldyrev, Reuven Cohen, Shlomo Havlin, and H. Eugene Stanley. Optimal paths in disordered complex networks. *Phys. Rev. Lett.*, 91(16):168701, October 2003. doi: 10.1103/PhysRevLett.91.168701.

CAIDA Macroscopic Topology Project Team. CAIDA skitter AS links topology (collection). http://imdc.datcat.org/collection/1-000W-X=CAIDA+skitter+AS+Links+Topology, 2000–2006.

Fan R. K. Chung. *Spectral Graph Theory*. Regional Conference Series in Mathematics. CBMS, 1997.

Bogdan Danila, Yong Yu, John A. Marsh, and Kevin E. Bassler. Optimal transport on complex networks. *Physical Review E*, 74(4):046106, 2006.

Holger Ebel, Lutz-Ingo Mielsch, and Stefan Bornholdt. Scale-free topology of e-mail networks. *Phys Rev E Stat Nonlin Soft Matter Phys*, 66(3 Pt 2A):035103, September 2002.

P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5:17–61, 1960.

Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 251–262, New York, NY, USA, August 1999. ACM Press. ISBN 1-58113-135-6. doi: 10.1145/316188.316229.

Linton C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, March 1977.

Henryk Fukś and Anna T. Lawniczak. Performance of data networks with random links. *Mathematics and Computers in Simulation*, 51(1-2):101–117, 1999. ISSN 0378-4754. doi: 10.1016/S0378-4754(99)00125-1.

Mark Granovetter. Social science. ignorance, knowledge, and outcomes in a small world. *Science*, 301(5634):773–774, August 2003. doi: 10.1126/science.1088508.

R. Guimerà, A. Díaz-Guilera, Fernando Vega-Redondo, A. Cabrales, and A. Arenas. Optimal network topologies for local search with congestion. *Phys. Rev. Lett.*, 89(24):248701, November 2002. doi: 10.1103/PhysRevLett.89.248701.

David G. Kendall. Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded markov chain. *The Annals of Mathematical Statistics*, 24(3):pp. 338–354, 1953. ISSN 00034851.

Wolfram Krause, Jan Scholz, and Martin Greiner. Optimized network structure and routing metric in wireless multihop ad hoc communication. *Physica A — Statistical Mechanics and its Applications*, 361(2):707–723, 2006. doi: 10.1016/j.physa.2005.06.085.

Dmitri Krioukov, Fan Chung, kc claffy, Marina Fomenkov, Alessandro Vespignani, and Walter Willinger. The workshop on internet topology (WIT) report. *ACM SIGCOMM Computer Communication Review*, 37(1):69–73, 2007. ISSN 0146-4833. doi: 10.1145/1198255.1198267.

Stanley Milgram. The small-world problem. *Psychology Today*, 1(1):61–67, May 1967.

Adilson E. Motter and Ying-Cheng Lai. Cascade-based attacks on complex networks. *Phys. Rev. E*, 66(6):065102, December 2002. doi: 10.1103/PhysRevE.66.065102.

Mark E. J. Newman. The structure of scientific collaboration networks. *Proc Natl Acad Sci U S A*, 98(2):404–409, January 2001. doi: 10.1073/pnas.021544898.

Mark E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(1):167–256, 2003.

Mirko Schäfer, Jan Scholz, and Martin Greiner. Proactive robustness control of heterogeneously loaded networks. *Physical Review Letters*, 96(10):108701, March 2006. doi: 10.1103/PhysRevLett.96.108701.

Jan Scholz, Wolfram Krause, and Martin Greiner. Decorrelation of networked communication flow via load-dependent routing weights. *Physica A — Statistical Mechanics and its Applications*, 387:2987–3000, January 2008. doi: 10.1016/j.physa.2008.01.011.

Bosiljka Tadić, G. J. Rodgers, and Stefan Thurner. Transport on complex networks: Flow, jamming and optimization. *International Journal of Bifurcation and Chaos*, 17(7):2363–2385, July 2007.

Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998. doi: 10.1038/30918.

Gang Yan, Tao Zhou, Bo Hu, Zhong-Qian Fu, and Bing-Hong Wang. Efficient routing on complex networks. *Physical Review E*, 73(4):046108, April 2006. doi: 10.1103/PhysRevE.73.046108.

Rui Yang, Wen-Xu Wang, Ying-Cheng Lai, and Guanrong Chen. Optimal weighting scheme for suppressing cascades and traffic congestion in complex networks. *Phys. Rev. E*, 79(2):026112, February 2009. doi: 10.1103/PhysRevE.79.026112.

Soon-Hyung Yook, Zoltán N. Oltvai, and Albert-László Barabási. Functional and topological characterization of protein interaction networks. *Proteomics*, 4(4):928–942, April 2004. doi: 10.1002/pmic.200300636.

L. Zhao, Y. C. Lai, K. Park, and N. Ye. Onset of traffic congestion in complex networks. *Physical Review E*, 71(2):026125, 2005.