

Optimization of transport protocols with path-length constraints in complex networks

José J. Ramasco*

Complex Systems Lagrange Laboratory (CNLL), ISI Foundation, 10133 Turin, Italy

Marta S. de La Lama†

*Max-Planck Institute for Dynamics and Self-Organization, Bunsenstr. 10, 37073 Göttingen, Germany;**Instituto de Física de Cantabria, CSIC-UC, Av. de los Castros s/n, 39005 Santander, Spain;**and Departamento de Física Moderna, Universidad de Cantabria, Santander, Spain*

Eduardo López‡

*CABDyN Complexity Centre, Saïd Business School, University of Oxford, Park End, Oxford OX1 1HP, United Kingdom**and Department of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom*

Stefan Boettcher§

Physics Department, Emory University, Atlanta, Georgia 30322, USA

(Received 3 June 2010; published 28 September 2010)

We propose a protocol optimization technique that is applicable to both weighted and unweighted graphs. Our aim is to explore by how much a small variation around the shortest-path or optimal-path protocols can enhance protocol performance. Such an optimization strategy can be necessary because even though some protocols can achieve very high traffic tolerance levels, this is commonly done by enlarging the path lengths, which may jeopardize scalability. We use ideas borrowed from extremal optimization to guide our algorithm, which proves to be an effective technique. Our method exploits the degeneracy of the paths or their close-weight alternatives, which significantly improves the scalability of the protocols in comparison to shortest-path or optimal-path protocols, keeping at the same time almost intact the length or weight of the paths. This characteristic ensures that the optimized routing protocols are composed of paths that are quick to traverse, avoiding negative effects in data communication due to path-length increases that can become specially relevant when information losses are present.

DOI: [10.1103/PhysRevE.82.036119](https://doi.org/10.1103/PhysRevE.82.036119)

PACS number(s): 89.75.Hc, 89.20.Hh, 05.70.Ln

I. INTRODUCTION

Communication systems are an essential element of the modern world, and represent some of the most noticeable settings of transport problems in complex structures. The Internet, telephone systems, power grids, etc. provide ways to transport information, signals, or resources from one part of the globe to another. Such communication is done by means of highly structured systems that can be modeled by networks, i.e., collections of elements corresponding to nodes, links that interconnect them, and weights on the links symbolizing connection capabilities [1–5].

In many examples of communication systems, discrete subsets of information (packets) are sent from a source to the appropriate destination. Typically, a multitude of users simultaneously perform similar tasks between many different source-destination pairs, and thus the network at any given time is populated by numerous traveling data packets. In order for packets to find their way, routing protocols are put in place to specify how the packets travel from source to destination. Optimization of the performance in communica-

tion systems can thus be achieved in two nonexclusive ways: improving the structure of transport networks [6] and finding optimal routing protocols [7]. Both approaches have been the focus of a great deal of research on this subject in recent years [6–25]. A number of important results have been obtained, elucidating the usefulness of the so-called scale-free networks [6,12], the vulnerability of networks to random or targeted failures of some of its elements [14], and the relevance of path lengths in the efficiency of communication [15].

When the overall traffic in the network is very large, the system may get compromised in its capability to deliver information. The origin of this problem could be twofold: either the connections of the network are unable to tolerate such a high traffic flow or, more commonly, some components (nodes) get overloaded, i.e., congested. Congestion occurs when the nodes become heavily used acting as bottlenecks for the network communication. By establishing clever routing protocols, this congestion can be diminished or, for moderate rates of overall traffic in the system, eliminated. The assessment of node workload requires thus the introduction of a new quantity. Betweenness centrality (betweenness in short) has been proposed as such a measure [6,17], as an estimate of how much use a given node can expect to receive. Specifically, for a set of paths in a network, the betweenness of a node is defined as the number of paths visiting that node. Therefore, a given routing protocol induces a betweenness on the nodes of the network. It has been pointed

*jramasco@isi.it

†msanche@gwdg.de

‡eduardo.lopez@physics.ox.ac.uk

§sboettc@emory.edu

out that nodes with large betweenness are the first to form packet queues when traffic is large enough [6]. If this queue is not eliminated, then the travel times of packets from source to destination increase without bound, hence, rendering the node (and potentially the network) congested.

On the other hand, packet delivery is an imperfect process since it can be subjected to packet loss or corruption. When congestion occurs, even entire queues of packets could be eliminated while waiting at overloaded nodes. Thus, when attempting to optimize a routing protocol, care must be taken not to increase the path length to a level that would elevate the likelihood of packet loss beyond critical limits. Such considerations have led to the formulation of models such as limited path percolation, which exhibits a phase transition when the length of paths increases beyond a specific tolerance level [15].

In this paper, we focus on the problem of optimizing the transport rules for scale-free networks with properties similar to those of computer networks. However, in contrast to a number of other studies, we optimize the protocols to reduce the likelihood of congestion *but* keeping path lengths as close as possible to the shortest, thus guaranteeing efficient communication paths. We find that by applying ideas borrowed from extremal optimization (EO) [26], which uses local information to approach the optimal state of a system, we can achieve both goals of congestion avoidance and short path lengths.

The paper is structured as follows. Section II explains the philosophy behind our optimization heuristic, the algorithm to find optimal routing protocols, and our network model. Section III quantifies betweenness of nodes as a consequence of the optimized routing protocol. Section IV focuses on the path lengths created by the EO-based routing protocol and the advantages that it affords regarding maintaining the path lengths within a constant factor of the original paths. Section V extends our methods to weighted networks and, finally, Sec. VI is dedicated to conclusions.

II. OPTIMIZATION METHOD

To define the betweenness precisely, we need first to specify a routing protocol, which is a set of paths $\mathcal{L}=\{\rho_{ij}\}$ between all pairs of nodes i and j of the network. We define the length of each path ℓ_{ij} as the minimum number of links needed to be crossed when going from i to j following the path ρ_{ij} . We adhere to the convention of considering only one path for each pair of nodes [20], as the length degeneracy of the paths will be exploited by the method to optimize the protocol. The rules by which these paths are chosen can be numerous. For instance, one typical choice for \mathcal{L} in the literature is the set of shortest paths (SPs) between nodes. The betweenness b_i of a node i is then defined as the number of paths of \mathcal{L} that visit node i , without counting those paths that begin or end in i . (This choice is somewhat arbitrary, and other choices are also possible.) Note that in the literature the word “betweenness” sometimes is associated exclusively with the number of paths crossing a node within the framework of the SP protocol. We are using this concept in a more general context since it can refer to any routing protocol,

which has to be given. In fact, betweenness values depend on both the structure of the network and \mathcal{L} . For a given network, changing the routing protocol typically changes b_i over the network except in very special cases such as when there is only one possible protocol choice (e.g., a star network).

Given a network and a protocol \mathcal{L} , it has been shown that the first node to begin accumulating a queue, say m , satisfies $b_m > b_i$ for all nodes $i \neq m$ in the network [6]. To explain this point, consider a network of size N in which data packets are produced in each node at random with a rate $0 \leq \gamma \leq 1$ per time step. In the simplest version of the problem, the destination of each packet is random too and, thus, each node can produce a packet to any of the other $N-1$ nodes of the network with equal probability. In this case, it was found that the number of packets visiting node i is, on average, $\gamma b_i / (N-1)$ [6]. The first node to congest, thus, is the one with the maximum betweenness value b_{max} [27]. We can also rephrase this point by indicating that the critical rate of packet production would be $\gamma_c = (N-1) / b_{max}$ before some nodes of the network congest [6]. This relation shows the relevance of designing routing protocols with a b_{max} as low as possible in order to prevent network congestion for low γ values. The lower b_{max} is, the higher the number of packets can be (higher γ) before the network congests. There is, however, a price to pay in this strategy because the reduction in b_{max} is usually attained at the expense of enlarging the path lengths. Under some circumstances, such as when packets can get lost, enlarging significantly the paths may not be the best option. The model that we propose next aims at bridging the gap between these two opposing forces: we explore how much b_{max} can be reduced while keeping the paths as close to their shortest version as possible.

We generate the networks for the analysis with the configuration model [28,29] that allows us to build graphs with arbitrary degree sequence but free from degree-degree correlations. The degree of a node i corresponds to its number of connections and is usually represented as k_i . In this work, we will focus on power-law degree distributed networks, $P(k) \sim k^{-\lambda}$, with exponent $\lambda=2.5$ to facilitate the comparison with previous studies [20,21,23]. Nevertheless, we have also performed the analysis on networks with $\lambda=4.5$ finding no qualitative difference in the final results.

The optimization algorithm starts from the SP protocol: once the network is built, the first step is to calculate the shortest paths between all nodes of the network using Dijkstra’s algorithm [30]. Then, in an iterative way for as long as desired (i.e., until an arbitrarily chosen stopping condition is reached [31]), we repeat the following steps:

- (i) The node with the highest betweenness m is selected.
- (ii) One of the paths, ρ_{ij} , passing through m is chosen at random; we refer to its initial and final nodes as i and j (with $i \neq m \neq j$).
- (iii) A path between i and j is searched. The new path must be the shortest alternative to ρ_{ij} not passing through m . In the case that there are more than one alternative path degenerate in length, one of them is selected at random. On the other hand, if there is no alternative to ρ_{ij} , the path remains without a change.

A schematic illustration of these rules is shown in Figs. 1(a) and 1(b). In the sketch, a path passing through the high-




FIG. 1. (Color online) Sketches (a) and (b) illustrate the way in which the method optimizes the protocol by selecting another path that must be either of the same length or the shortest alternative to the original one. In (c), the probability of finding alternative length-degenerate paths, $P_{n>1}$, for the shortest-path protocol in networks with $\lambda=2.5$ is displayed as a function of the network size N . Inset: histogram for the number n of degenerate alternatives for each path in the SP protocol. The different curves correspond to the following network sizes: $N=100$ (circles), 500 (squares), 2000 (diamonds), and 10 000 (triangles). The x-axis is log-binned to facilitate visualization.

est betweenness node, the square, is rerouted thereby reducing betweenness. It is worth stressing that the constraint on the length of the alternative paths in practice means that if another choice of equal length to the original path exists, the path does not increase in length. The method takes advantage of the possible degeneracy of the paths with respect to the length to change the protocol and to ease the stress on the node(s) with the largest traffic. Note also that the method keeps no long-term memory; the search for alternative paths due to sequential expulsions of a path from different nodes does not accumulate. The algorithm generates thus protocols with paths limited to SP configurations or the shortest alternatives to them avoiding a single high used node. For this reason, we will call the protocols resulting from the method application optimized shortest paths (OSPs) from now on, a term that we will also use, by extension, to refer to the method itself.

The potential utility of exploiting the SP degeneracy and its importance in our networks can be seen in Fig. 1. We have plotted in Fig. 1(c) the probability of finding a degenerate alternative to the shortest path between any pair of nodes i and j , $P_{n>1}(N)$, as a function of the network size N , where n is the number of paths with equal (shortest) length between a pair of nodes. $P_{n>1}(N)$ consistently grows with N , becoming higher than 65% for the largest graphs that we have considered, $N=10\,000$. In the inset, the distribution of the number of different degenerate paths, $P(n)$, is also shown for four network sizes. Note that $P_{n>1} = \sum_{n=2}^{\infty} P(n)$. These plots correspond to networks with $\lambda=2.5$, as mentioned above, changing λ alters the values of the exponents found in these curves [inset in Fig. 1(c)] but not the main qualitative picture.




FIG. 2. (Color online) Betweenness distribution for a $N=500$ system. The (black) circles correspond to the stationary distribution reached with our optimization method, while the (red) squares are for the initial SP protocol. Inset: the stationary betweenness distributions of the method for other two network sizes: $N=350$ and $N=700$.

The philosophy of OSP is in the spirit of EO [26], attacking the worst element of the system in an attempt to improve the global performance. In the following, we will study how its application affects the different features of the protocols.

III. BETWEENNESS DISTRIBUTION

We will discuss first the results concerning the betweenness of the nodes and how its values adapt after the optimization method is employed. After some iterations of the method [31] the betweenness distribution $P(b)$ reaches a stationary form that can be seen in Fig. 2. Before describing it, it is important to remember that the SP protocol produces a power-law distribution for $P(b)$ whose exponent depends on the exponent λ of the degree distribution [17,18] [see the curve of squares (red online) in Fig. 2]. The protocols generated by the optimization method show that the tail of $P(b)$ collapses into a peak at high values of the betweenness. Also the area of low b suffers a slight variation, although the functional form of $P(b)$ at intermediate b seems largely unaffected. A change in the size of the network displaces the position of the peak but not the quality of the effects observed in the main plot of Fig. 2.

The peak induced in $P(b)$ by the optimization method at high values of b correlates with a decline of the maximum betweenness in the network. As mentioned before, this betweenness, b_{max} , is of crucial importance for the transport capability of a protocol since it imposes an upper cutoff to the traffic that the network is able to sustain before congesting [6,13]. The scaling of b_{max} with the network size N can thus be seen as a measure of the scalability of a routing protocol: the faster b_{max} grows with N , the less useful the protocol is for large networks. In Fig. 3, the average b_{max} in the stationary state of the protocols obtained with the OSP method is depicted as a function of the number of nodes in the network. We find a power-law increase with a functional form $b_{max} \sim N^\alpha$ and $\alpha^{OSP} \approx 1.3$. This value of α must be compared with the results encountered with other protocols or other protocol optimization methods. In the same figure, we have also included a table with the exponents α for other




FIG. 3. (Color online) The average maximum betweenness b_{max} as a function of the network size N for the optimized protocols. Lines are shown as a reference scaling with exponents 1.3 (solid line), 1.18 (dashed line), and 1.62 (dotted line). In the table, scaling exponents α obtained with different optimization methods are listed. The initials correspond to shortest path (SP) [32], hub avoidance (HA) [20], Danila’s method (DAN) [21], efficient routing (ER) [23], and the OSP method.

methods. It is specially interesting to compare with the original shortest-path protocol $\alpha^{SP} \approx 1.62$ [32] or with the best of the optimization methods listed (Danila’s method) with $\alpha^{DAN} \approx 1.18$ [21,22]. OSP, without deviating substantially from the SP protocol in the length of the paths, produces an acceptable value of the exponent α .

IV. PATH LENGTHS

The near invariance of the path length between SP and our optimized protocol, and the path-length distribution itself, deserves further attention. Here, we will focus on studying how the length of the paths in the protocol mutates when the optimization method is applied to the protocol. In Fig. 4, we have plotted the path-length distribution $P(\ell)$ for the original shortest-path protocol (red squares) and for the stationary regime of the protocols obtained with the OSP method on a network of size $N=1000$. As can be seen, the distribution is slightly displaced to larger values of ℓ , but its shape remains essentially invariant. In order to quantify the relevance of the change of $P(\ell)$ with the system size, we have represented in




FIG. 4. (Color online) (a) Path-length distribution for a network of size $N=1000$. In the figure, the results of the protocols obtained with the OSP method can be compared with the initial SP. In (b), the ratio between the average lengths of the paths in the two protocols (OSP and SP) is visualized as a function of the network size.




FIG. 5. Average path length $\langle \ell \rangle$ of the protocols obtained as a function of the network size. It can be seen that a $\log[\log(\cdot)]$ scale in (b) provides a better fit to our simulation results than a pure logarithmic scale in (a).

Fig. 4(b) the ratio between the average lengths of the paths in the protocols obtained with OSP and that calculated with the SP. The curve is monotonically decreasing with N , moving very slowly toward unity (or a value just above).

Similar to b_{max} , limiting the scaling of the average length of the paths, $\langle \ell \rangle$, with the system size is also an important design feature. A protocol that elongates the paths unnecessarily may not be efficient for network communications. Generally, there exists a nonzero probability of losing a packet in every communication between two servers, and the longer the paths become, the higher the probability of missing information. Furthermore, if data loss reaches a point in which most of the paths are not functioning, the network suffers a disruption similar to a percolation transition. In fact, this phenomenon has been studied recently, and it is known as limited path percolation [15]. In Fig. 5, we show how $\langle \ell \rangle$ in the protocols obtained using our OSP method behaves with increasing network sizes. In scale-free networks with $\lambda=2.5$, the SP protocol $\langle \ell_{SP} \rangle$ is expected to grow as $\log \log(N)$ [33]. We cannot numerically test this formula for multiple decades in N for OSP protocols, but within the limited range of values that we can explore the relation holds. Such scaling is important because, as mentioned above, the optimization of the protocols can in some cases imply a trade-off between reductions in b_{max} and escalating path lengths. For instance, Danila’s method [21], which exhibits the best scaling for b_{max} versus N of the methods shown in Fig. 3, leads to a $\langle \ell \rangle \sim \log(N)$ [22], considerably faster than the $\log[\log(N)]$ scaling of the OSP method (or of the SP protocol). This difference in the scaling can be relevant if the packet communication has losses in long paths (or in paths longer than the SP).

In general, consider a protocol that is tuned to scale with the typical structural path lengths of the network, which—to be concrete—we associate with the most probable length ℓ_{SP}^* of the length distribution $P_{SP}(\ell)$. We imagine that the real system uses a protocol \mathcal{L}_o when uncongested which is, for our purposes, the best practical protocol for the network of interest (even if it is not strictly optimal). For a power-law network with exponent $\lambda=2.5$, $\ell_o^* = D_o \ell_{SP}^* \sim D_o \log \log(N)$, where $D_o \geq 1$ is a numerical constant. Now, when higher rates of packet creation appear on the network, to avoid congestion we consider introducing two new protocols \mathcal{L}_S and \mathcal{L}_L , the first one with typical path lengths $\ell_S^* \sim \delta_S \ell_o^*$ ($\delta_S > 1$ another constant) and the second with $\ell_L^* = F(\ell_o^*)$, with $F(x)$




FIG. 6. (Color online) Cumulative distribution of the ratio between the lengths of the paths obtained with our protocol optimization method and those of the corresponding original SP. The distribution becomes steplike due to the discreteness of the length.

$> x^{1+\epsilon}$, where $\epsilon > 0$. In the thermodynamic limit, with mean-field approximation, τ , the ratio of new typical path length to old typical path length in \mathcal{L}_S , remains constant as we can see from

$$\tau_S = \frac{\ell_S^*}{\ell_o^*} = \frac{\delta_S D_o \log \log(N)}{D_o \log \log(N)} \rightarrow \delta_S. \quad (1)$$

On the other hand, for \mathcal{L}_L , we find

$$\tau_L = \frac{F[D_o \log \log(N)]}{D_o \log \log(N)} \rightarrow [D_o \log \log(N)]^\epsilon \rightarrow \infty, \quad (2)$$

where the last relation occurs for $N \rightarrow \infty$. In Ref. [15], the limited path percolation transition was shown to be intimately related to τ . For our purposes, it means that if a protocol keeps τ finite as $N \rightarrow \infty$, occurring in the τ_S case, there is the possibility to have a protocol that can be tuned to maintain at least a level of communication between node pairs. Note that the \mathcal{L}_S protocol can be written as $F(x)=x$. On the other hand, if τ diverges, as in τ_L , as the system size grows, the typical pair of nodes becomes increasingly distant in the relative sense, and thus cannot in general remain in communication; or, in a more formal sense, protocols such as \mathcal{L}_L may not be scalable [34].

The protocol we propose here has the behavior of \mathcal{L}_S and, thus, as in limited path percolation, can have a percolation threshold to a connected regime. Furthermore, the fact that we study scale-free degree distributed networks with small λ indicates that from the standpoint of limited path percolation, there is always a connected regime. In addition, when δ is close to 1, the protocol is very robust as most pairs of nodes remain connected after the introduction of the optimization of b_{max} . These conclusions are also supported by results displayed in Figs. 4–6.

Furthermore, as explained above, a path obtained with our method can be only in two states: either in the SP configuration or, if they were expelled from a high used node and no length-degenerate alternative was available, in the shortest configuration that does not pass through that node. Such a two-state configuration recalls physical systems in which the components can be in the energetic ground state or being

activated by thermal (stochastic) fluctuations to higher energy levels. To confirm whether this parallelism is valid, we have plotted in Fig. 6 the cumulative distribution [$C_>(R) = \int_R^\infty dR' P(R')$] of the ratios $R = \ell_{stat} / \ell_{SP}$, where ℓ_{stat} corresponds to the length of each path in an stationary instance of the protocols obtained with our optimization technique and ℓ_{SP} is the corresponding length of the shortest path. As can be seen, the probability for the paths to be outside the ground state (SP), in an “excited state,” decays as an exponential function with a characteristic ratio that depends on the graph (whether the graph is scale-free or not and, if it is so, also on the exponent of the degree distribution). In the example of Fig. 6, the characteristic ratio is around 1.55, which makes it extremely unlikely for any path to suffer a large stretching out of its SP length. This exponentially decaying $C_>(R)$ resembles the Boltzmann factor in which $R-1$ plays the role of the energy.

V. WEIGHTED NETWORKS

The introduction of disorder in the network connections deeply modifies the previous scenario. However, connections of diverse quality can be expected in many real-world situations. A classical example within the context of transport is the Internet and the variability on bandwidth of the connections between servers. Mathematically speaking, the quality of a connection can be represented by a scalar attached to each link that is known as its weight w . The presence of weighted links alters the definition of shortest path [35–37]. Instead of “short” in the sense of number of links between origin and destination, it may be important to find the path with the lowest weight along the way (or the largest, depending on the quality that the weight conveys). These paths are often referred to as optimal paths (OPs) in the literature (see, for instance, [38]).

Protocols based on OPs also have a very particular scaling with the network size [13]. Depending on the level of randomness on the weights of the connections, the system can fall into a strong- or a weak-disorder regime. The difference between these two regimes is that in the strong-disorder case the fluctuations on the weight are so large that the accumulated weight along each path is controlled by the edge with the highest weight, while in weak-disorder the “responsibility” for the path’s overall weight is distributed among all the links along the path [39,40]. In the two regimes, the scaling of b_{max} and of the average weight of the paths, $\langle w_{path} \rangle$, with N changes with respect to unweighted graphs [25]. Here, we will focus only on the weak-disorder regime since it is the only one in which it makes sense to search for an alternative protocol to the OP. The cost of doing so in the strong-disorder limit would diverge with the size of the graph.

One of the positive aspects of the OSP protocol optimization is that its generalization to weighted graphs is straightforward: at each step, a path passing through the node with the highest betweenness is chosen to be redirected. The path is then exchanged by its alternative of lowest cost that does not pass through that node. As before, if there is no alternative, the path remains invariant. In the following, we will study the performance of the stationary protocols produced




FIG. 7. (Color online) Dependence of b_{max} with the system size for weighted graphs. The (red) circles are the results for the protocols of OSP, and the (blue) squares are for the original OP protocol. The weight distribution in (a) follows an exponential $P(w) \sim e^{-w/2}$, and in (b) the power law $P(w) \sim w^{-2.5}$, as displayed in the insets for networks of size $N=500$. The continuous curves in both panels are a guide to the eye showing a visual comparison between the exponential values obtained by fitting the simulation results.

by this method on Molloy-Reed graphs with $\lambda=2.5$ and with two possible functional forms for the weight distributions: either an exponential, $P(w) \sim e^{-w/w_c}$, or a power law, $P(w) \sim w^{-\beta}$. In order to keep the results comparable for both distributions, we fix the parameter w_c to obtain the same average weight as in the power-law distribution. For the data shown, this will be $w_c=2$ and $\beta=2.5$.

The scaling of b_{max} as a function of the network size is shown in Fig. 7 for networks with both types of weight distributions. The application of the optimization method produces a significant improvement in b_{max} with respect to the optimal-path protocol. For the range of system sizes of the figure, this entails an improvement by a factor of 3 or even higher. Additionally, the exponent α of the optimized protocols is close to 1.4 while that of the original OP protocols was about 1.55. One can wonder about the price to pay for such an improvement. In networks with continuous values of the weights, any degeneracy in the weight of the alternative paths is extremely unlikely. Strictly speaking, then, there is always a price to pay, although it could be minute. To answer




FIG. 8. (Color online) Comparison between the average weights of the paths, $\langle w_{path} \rangle$, in the original OP protocol and in the optimized protocols with OSP. The curves in (a) correspond to weighted graphs with an exponentially decaying weight distribution, while those in (b) have a power-law weight distribution.

this question, we have plotted in Fig. 8 the average weight of the paths scaled with N for protocols obtained with our method and for the corresponding OP on both types of networks and both weight distributions. Similarly to the unweighted case, we find that the average weight is slightly higher, but not extremely so due to being in the weak disorder regime. Also the ratio between the average weights of the paths of our OSP optimized protocols and of OP decreases for increasing network sizes. Furthermore, even though it is not shown in the figure, neither the shape of the distribution nor the weight of the paths is hardly altered by the optimization of the protocols.

VI. DISCUSSION AND CONCLUSION

We study a routing protocol optimization method that takes advantage of the degeneracy of the shortest paths to reduce the congestion in the nodes of the network. The paths between node pairs in the optimized protocols can only be in two states: they can be either one of the shortest paths or the closest alternative avoiding highly used nodes. We use ideas coming from extremal optimization to guide our algorithm, which proves to be an effective technique. The resulting routing protocols show an appreciable reduction in the maximum node betweenness, b_{max} , while keeping the path-length distribution similar to the lower bound provided by the SP protocol. For the networks that we have analyzed, b_{max} increases with N as a power law with an exponent $\alpha \approx 1.3$. That is only slightly higher than the best values reported in the literature, and—in contrast to those methods—the average path length appears to retain the $\log[\log(N)]$ scaling of the SP protocol. We show that methods which do not conserve a constant ratio of path length with the topological shortest paths end up in a disconnected phase of percolation with limited paths, which is a possible indicator that on networks with packet loss such methods may not be scalable.

Our optimization method, OSP, is easily generalized for weighted networks. We have studied two examples of weighted networks in the weak-disorder regime and found significant reductions in b_{max} with respect to the optimal-path protocol as well as in the scaling exponent with which b_{max} grows with N .

As we implemented the optimization method—a procedure in which we keep record of all the paths passing through each node—the numerical performance is not optimal. Indeed, we were able to study a range of network sizes wide enough to characterize the scaling of the optimized protocols properties, but not as large as to allow for large network analysis. Our aim in this work was to explore the behavior of the optimization method and to ascertain whether the resulting routing protocols have interesting features concerning the size scaling and behavior of b_{max} and the path-length distributions. Indeed, the simple fact that we have shown that it is feasible to perform such dual optimization is an important step forward. Furthermore, we have found that this two-parameter optimization is important because without taking into account both b_{max} and path length, the protocols that simply reduce b_{max} may be less scalable.

Therefore, the possibility of searching for efficient numerical methods to implement the optimization algorithm remains open for future work. A possible solution, depending on the structure of the network, can be the implementation of a local search technique for the alternative paths. This local method could be highly efficient if the clustering (the density of triangles) in the graph is high as typically occurs in geographically embedded or social networks. Locally treelike structured networks, such as those considered here, typically require *global* rearrangements between paths of nearly degenerate length but hardly any overlap, as illustrated in Figs. 1(a) and 1(b). Maintaining the paths of the routing protocol short can be crucial in situations such as when the system

suffers packet losses, a circumstance that seems to be relatively frequent in practice. OSP represents in such cases a useful tool that offers a balance between betweenness reduction (lower congestion) and restricted path lengths (reduced losses) in routing protocols.

ACKNOWLEDGMENTS

E.L. acknowledges funding from the EPSRC Grant No. EPSRC EP/E056997/1. J.J.R. is funded by the FET Grant No. 233847-Dynanets of the EU Commission. Also both J.J.R. and S.B. acknowledge support by the U.S. National Science Foundation through Grant No. DMR-0312510.

-
- [1] R. Albert and A.-L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
 [2] M. E. J. Newman, *SIAM Rev.* **45**, 167 (2003).
 [3] S. Dorogovtsev and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford University Press, New York, 2003).
 [4] R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge University Press, Cambridge, England, 2004).
 [5] A. Barrat, M. Barthélemy, and A. Vespignani, *Dynamical Processes on Complex Networks* (Cambridge University Press, Cambridge, England, 2008).
 [6] R. Guimerà, A. Díaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas, *Phys. Rev. Lett.* **89**, 248701 (2002).
 [7] T. Ohira and R. Sawatari, *Phys. Rev. E* **58**, 193 (1998).
 [8] A. Arenas, A. Díaz-Guilera, and R. Guimerà, *Phys. Rev. Lett.* **86**, 3196 (2001).
 [9] R. V. Solé and S. Valverde, *Physica A* **289**, 595 (2001).
 [10] P. Echenique, J. Gómez-Gardeñes, and Y. Moreno, *Phys. Rev. E* **70**, 056105 (2004).
 [11] P. Echenique, J. Gómez-Gardeñes, and Y. Moreno, *EPL* **71**, 325 (2005).
 [12] E. López, S. Carmi, S. Havlin, S. V. Buldyrev, and H. E. Stanley, *Physica D* **224**, 69 (2006).
 [13] L. A. Braunstein *et al.*, *Int. J. Bifurcation Chaos* **17**, 2215 (2007).
 [14] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, *Phys. Rev. Lett.* **86**, 3682 (2001).
 [15] E. López, R. Parshani, R. Cohen, S. Carmi, and S. Havlin, *Phys. Rev. Lett.* **99**, 188701 (2007).
 [16] L. A. Braunstein, S. V. Buldyrev, R. Cohen, S. Havlin, and H. E. Stanley, *Phys. Rev. Lett.* **91**, 168701 (2003).
 [17] K.-I. Goh, B. Kahng, and D. Kim, *Phys. Rev. Lett.* **87**, 278701 (2001).
 [18] M. Barthélemy, *Phys. Rev. Lett.* **91**, 189803 (2003).
 [19] S. Sreenivasan, T. Kalisky, L. A. Braunstein, S. V. Buldyrev, S. Havlin, and H. E. Stanley, *Phys. Rev. E* **70**, 046133 (2004).
 [20] S. Sreenivasan, R. Cohen, E. López, Z. Toroczkai, and H. E. Stanley, *Phys. Rev. E* **75**, 036105 (2007).
 [21] B. Danila, Y. Yu, J. A. Marsh, and K. E. Bassler, *Phys. Rev. E* **74**, 046106 (2006).
 [22] B. Danila, Y. Yu, J. A. Marsh, and K. E. Bassler, *Chaos* **17**, 026102 (2007).
 [23] G. Yan, T. Zhou, B. Hu, Z.-Q. Fu, and B. H. Wang, *Phys. Rev. E* **73**, 046108 (2006).
 [24] T. Kalisky, L. A. Braunstein, S. V. Buldyrev, S. Havlin, and H. E. Stanley, *Phys. Rev. E* **72**, 025102(R) (2005).
 [25] K.-I. Goh, J. D. Noh, B. Kahng, and D. Kim, *Phys. Rev. E* **72**, 017102 (2005).
 [26] S. Boettcher and A. G. Percus, *Phys. Rev. Lett.* **86**, 5211 (2001); *Artif. Intell.* **119**, 275 (2000).
 [27] We focus on the limiting case of the network in which all nodes have equal capacity to process packets per unit time (which we take to be unity for simplicity). This choice highlights, in a sense, the worst case scenario where nodes are considered regardless of their other characteristics such as degree. In practice, the situation is likely to be different as, e.g., highly connected nodes would be expected to have extra packet processing capacity.
 [28] M. Molloy and B. Reed, *Random Struct. Algorithms* **6**, 161 (1995).
 [29] M. Catanzaro, M. Boguñá, and R. Pastor-Satorras, *Phys. Rev. E* **71**, 027103 (2005).
 [30] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms* (McGraw-Hill, New York, 2005).
 [31] As a criterion to define the stationarity of the optimization process, we focus on the value of b_{max} as a function of the iteration steps. b_{max} is considered over the last 20 steps, and the process is said to be stationary if the relative error $E(b_{max}) \leq 10^{-4}$. We confirmed the stationarity of the betweenness and path-length distributions within numerical precision once this criterion is fulfilled. The number of path rearrangements required to achieve stationary protocols scales for the networks studied approximately as $\sim N^{2.5}$.
 [32] Y. Yu, B. Danila, J. A. Marsh, and K. E. Bassler, *EPL* **79**, 48004 (2007).
 [33] R. Cohen and S. Havlin, *Phys. Rev. Lett.* **90**, 058701 (2003).
 [34] Note that any function $F(x)$ scaling faster than x and diverging is sufficient to obtain the same qualitative results. For instance, $F(x)=x \log x$ has the same effect. In general, if $\lim_{x \rightarrow \infty} F(x)/x \rightarrow \infty$ the paths scale too fast for a protocol satisfying limited paths to be scalable.
 [35] K. Park, Y.-C. Lai, and N. Ye, *Phys. Rev. E* **70**, 026109 (2004).
 [36] M. E. J. Newman, *Phys. Rev. E* **64**, 016132 (2001).

- [37] M. E. J. Newman and M. Girvan, *Phys. Rev. E* **69**, 026113 (2004).
- [38] M. Cieplak, A. Maritan, and J. R. Banavar, *Phys. Rev. Lett.* **72**, 2320 (1994).
- [39] Y. Chen, E. López, S. Havlin, and H. E. Stanley, *Phys. Rev. Lett.* **96**, 068702 (2006).
- [40] Z. Wu, E. López, S. V. Buldyrev, L. A. Braunstein, S. Havlin, and H. E. Stanley, *Phys. Rev. E* **71**, 045101(R) (2005).