

Using relaxational dynamics to reduce network congestion

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2008 New J. Phys. 10 093007

(<http://iopscience.iop.org/1367-2630/10/9/093007>)

[The Table of Contents](#) and [more related content](#) is available

Download details:

IP Address: 128.197.42.52

The article was downloaded on 01/04/2010 at 13:41

Please note that [terms and conditions apply](#).

Using relaxational dynamics to reduce network congestion

Ana L Pastore y Piontti^{1,6}, Cristian E La Rocca¹,
Zoltán Toroczkai^{2,3}, Lidia A Braunstein^{1,4}, Pablo A Macri¹
and Eduardo López^{3,5}

¹ Instituto de Investigaciones Físicas de Mar del Plata (IFIMAR)—
Departamento de Física, Facultad de Ciencias Exactas y Naturales,
Universidad Nacional de Mar del Plata—CONICET, Funes 3350,
7600 Mar del Plata, Argentina

² Department of Physics and Center for Complex Network Research,
University of Notre Dame, 225 Nieuwland Science Hall, Notre Dame, IN,
46556, USA

³ Theoretical Division, Los Alamos National Laboratory, Mail Stop B258,
Los Alamos, NM 87545, USA

⁴ Center for polymer studies, Boston University, Boston, MA 02215, USA

⁵ CABDyN Complexity Cluster and Department of Physics,
University of Oxford, Park End Street, Oxford OX1 2DF, UK

E-mail: apastore@mdp.edu.ar

New Journal of Physics **10** (2008) 093007 (9pp)

Received 14 May 2008

Published 5 September 2008

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/10/9/093007

Abstract. We study the effects of relaxational dynamics on congestion pressure in scale-free (SF) networks by analyzing the properties of the corresponding gradient networks (Toroczkai and Bassler 2004 *Nature* **428** 716). Using the Family model (Family and Bassler 1986 *J. Phys. A: Math. Gen.* **19** L441) from surface-growth physics as single-step load-balancing dynamics, we show that the congestion pressure considerably drops on SF networks when compared with the same dynamics on random graphs. This is due to a structural transition of the corresponding gradient network clusters, which self-organize so as to reduce the congestion pressure. This reduction is enhanced when lowering the value of the connectivity exponent λ towards 2.

⁶ Author to whom any correspondence should be addressed.

Transport networks, such as computer networks (Internet), airways, energy transportation networks, etc are amongst the most vital components of modern day infrastructures. These large-scale networks have not been globally designed, instead they are the result of local processes. It has been observed that many networks have a scale-free (SF) connectivity structure [1, 2]. SF networks are characterized by a power-law degree distribution $P(k) \sim k^{-\lambda}$, ($k \geq k_{\min}$), where λ is the connectivity exponent and k_{\min} is the lower degree that a node can have. There have been a number of mechanisms proposed in the form of stochastic network growth models that produce SF structures [1, 2] including weighted versions of these processes [3] (for discussions on the utility of SF models see [4]). However, these models do not *explicitly* connect the flow dynamics and transport performance (such as throughput, queuing characteristics, etc) with network topology. This is difficult to do in general, since the timescales of the flow on the network and that of the network's structural evolution itself can be rather different.

In this paper, we show that the emergence of SF structures is favored against non-SF structures, such as random graphs, if the transport dynamics has a relaxation component (called load-balancing in communications). We will see that *even one step* of such a gradient flow [5] will considerably reduce the congestion pressure in SF networks, whereas it has virtually *no effect* in random graphs. In addition, within the class of uncorrelated SF networks, the congestion reduction is enhanced for low (close to 2) λ values. Although we use the jargon from the fields of communication networks and queuing theory, we expect that our results hold for large-scale networks in general, where the flow dynamics is induced at least in part by the existence of gradients, a rather ubiquitous mechanism. In the following, by 'packet' we mean any discrete entity transported between two nodes of a network of N nodes. We assume that the network is driven in 'the volume', by packets entering at random (following a Poisson process) at a rate γ (the expected number of arrivals per unit time) at any of the nodes [6]–[9]. This is realistic since the users' actions in general are uncorrelated. Using the language of queuing theory, if a node in the time interval $(t, t + \tau)$ sends packets into the network, but it receives no packets from any of its neighbors, we say that it acts as a 'client', while if it receives a packet or several more from its neighbors, we say that it acts as a 'server'⁷. Here, $\tau > 0$ is the average processing time of a single packet by a node. Measuring the average fraction $J = \langle N_c/N \rangle$ of the number of clients N_c (over a period of time in the steady state), gives us a simple global measure for the congestion *pressure* present in the network [5]. Note that this is equivalent to $J = 1 - \langle N_{\text{servers}}/N \rangle$ [5, 10]. From this definition, higher J means more congestion (fewer servers handling the total network flow). The average $\langle \cdot \rangle$ is over network structure but it can also be over the randomness in the input when comparing *classes* of networks. The client nodes are the ones that introduce new packets into the network, but they do not contribute to routing. Certainly, all networks will become congested at large enough driving rates $\gamma > \gamma_c$ [6]–[9], [11, 12]. J indicates which network will become congested earlier, larger J meaning smaller γ_c . J is a global indicator that, however, does not take into account the *distribution* of the packets over the server nodes. That can be done via betweenness-based quantities as in [13]. Load-balancing is a specific case of the more general process of gradient-induced flows [5], where the flows are produced by the local gradients of a non-degenerate scalar field $\mathbf{h} = \{h_i\}_{i=1}^N$ distributed over the N nodes of a substrate graph G (transport network). The scalar field could represent, for example, the number of packets at the routers [14, 15], or the virtual time horizon of the

⁷ Packet flow is a continuous-time discrete event process. However, in the non-infinitesimal time $\tau > 0$ of processing a single packet, several packets may enter a node's queue from its neighbors or from being generated at that node.




Figure 1. Congestion pressure in the static model J_S and after relaxation J_R as function of N for SF networks with different values of λ , 2.5 (\circ), 3 (\diamond), 3.5 (\square), with filled symbols for S and empty symbols for R. In the inset, we plot J_S , (\circ), and J_R , (\square), as function of N for ER networks (here $p = 0.1$, however, other values led to similar results). The results were averaged over 10^3 network realizations.

processors in parallel discrete event simulations [16]. The gradient direction of a node i is a directed edge pointing towards neighbor j (on G), of i which has the lowest value of the scalar in its neighborhood. If i has the lowest value of h in its network neighborhood, the gradient link is a self-loop. The gradient network $\nabla_h G$ is defined simply as the collection of all gradient edges on the substrate graph G [5, 10]. It represents the subgraph of *instantaneous maximum flow* if the flow is induced by these gradients. In the gradient network, each node has a unique outgoing link and ℓ incoming links. When a node has $\ell = 0$ (no inflow in that instant), it acts as a client, otherwise it is a server. Then certainly, J is the average fraction of nodes, with $\ell = 0$, i.e. it is the fraction of the ‘leaves’ of $\nabla_h G$. Note that J is a queuing characteristic, rather than an actual throughput measure. It was shown [5, 10] that distributing random scalars $\{h\}$ independently onto the nodes of a network G , to which we refer as the *static* (S) model in the remainder, Erdős–Rényi (ER) graphs [17] (with given link probability p)⁸ become more congested with increasing network size N , i.e. $J_S \rightarrow 1$ while on SF networks J_S converges to a finite sub-unitary value, see the plots for J_S in figure 1. Gulbahce [18] introduced a Monte-Carlo update scheme adjusting the potential of each node by small random (Gaussian distributed) amounts and accepting only those changes that lower the congestion pressure. By *construction*, this procedure decreases the jamming in both ER and SF networks, with significant decrease observed for SF networks with $\langle k \rangle > 2$. As we will see below, a simple relaxational dynamics, in which the packet flow is driven by local gradients (i.e. load-balancing, a natural process) will also reduce drastically the jamming, especially in the case of SF networks. This observation, together with [18], strengthens the belief that SF networks are better suited to reduce congestion under gradient-driven flows.

⁸ We present our simulations on ER for $p = 0.1$, but we check that the result is valid for other values as well.

Usually, packets have destinations, and thus they cannot be governed exclusively by gradient flows, however, relaxational dynamics can be employed for finite periods of time. Once the (gradient) flow commences through the network, the scalar field becomes correlated and the queuing characteristics change. In the following, we systematically study the effects of a *single* relaxational step and show that, even in this case, the effects on congestion pressure can be drastic as the network becomes more heterogeneous. First, we note that the one-step relaxation dynamics defined by the gradient flow is nothing but the deposition model with surface relaxation (Family model) [19] from surface-growth physics extended to networks [20].

To generate SF networks, we used the configurational model [21] with $k_{\min} = 2$ (mostly for mathematical convenience, but see the discussion at the end about other networks)⁹. At $t = 0$, a random scalar field h is constructed by assigning to each node of the substrate network a random scalar independently and uniformly distributed between 0 and 1. At this stage the initial static gradient network [5] is formed and its jamming coefficient J_S determined. Then the scalars $h \equiv h(t)$ are evolved obeying the rules of the Family model [19]: at every time step a node i of the substrate is chosen at random with probability $1/N$ and it becomes a candidate for growth. If $h_i < h_j$, for every j (gradient criterion) which is a nearest neighbor of the node i , $h_i \rightarrow h_i + 1$. Otherwise, if h_i is not a minimum, the node j with minimum h is incremented by one. When the process reaches the steady state [20] of the evolution with this relaxation (R), we construct the gradient network and measure the pressure congestion J_R after the relaxational dynamic is performed. In accordance with previous observations [20, 22], the steady state is reached extremely quickly: the saturation time actually does not scale with the system size N but it approaches an N -independent constant [23, 24].

In figure 1, we plot J_S and J_R as function of N for ER graphs and for SF networks (for the latter, we compare cases with different λ values). As one can observe (inset), for ER networks, the model with relaxation has no effect on lowering the congestion, i.e. $J_S \simeq J_R$. For SF networks, however, there is a drastic difference between the static and dynamic cases, with J_S being considerably larger than J_R for large enough N . Note that J_S increases with decreasing λ , which can be understood through the fact that for lower λ values the $\nabla_h G$ of the SF graph is increasingly star-like, creating more congestion (see below). Since in real-world networks, however, one expects to find a load-balancing component of the transmission dynamics (see [14]), our model with relaxation is a better representation than the static one. And indeed, from figure 1 it becomes apparent that J_R has the *opposite* behavior as function of λ for large enough networks: lowering λ lowers the congestion pressure J_R . Next we show how the drop in the congestion pressure due to relaxation dynamics can be understood in terms of a structural change of the clusters of the corresponding gradient network. Here, the clusters are defined as the disconnected components (trees) of the gradient graph. The decrease of J_R (which is the fraction of leaves of the gradient network) means a decrease of the perimeter of the clusters of $\nabla_h G$.

In figure 2, we show J_S and J_R as function of λ for a fixed network size N . From this we can see that for a given value of N , J_S is larger than J_R . This is compatible with a transition on the structure of the trees of the gradient network from a star-like structure in the S state to a more elongated structure in the R (see figures 3(a)–(d)). This transition is responsible for the drop in the congestion pressure after the relaxation step is applied. From figure 3 we can see that for the ER networks the R model does not significantly affect the structure of the clusters

⁹ For $k_{\min} > 1$ there is a high probability that the network is fully connected, which is required in our model, see also [20].




Figure 2. Plot of the congestion pressure J_R , \circ and J_S , \square , for SF networks as function of λ ($N = 4000$). The results were averaged over 10^3 network configurations.




Figure 3. Largest clusters of the gradient networks of ER and SF with $\lambda = 2.5$ and $N = 1024$. In the top plots, we show the structure for the ER and on the bottom for SF. The left plots ((a) and (c)) correspond to the S and the right ones ((b) and (d)) to the R models.

of the gradient network (going from figure 3(a) to (b)), meanwhile, for SF networks (going from figure 3(c) to (d)) we observe a major structural transition before and after the relaxation dynamics is applied, from a star-like cluster to an elongated one.

For a quantitative insight into the S model consider that the scalars are identically (and independently) distributed according to some distribution $\eta(h)$. For the calculations below, we assume that the SF network has low clustering, i.e. that the probability of two neighbors of a node to be also neighbors is small. Assume that the central node i has its scalar value equal to h_i . Then, the probability that the *neighbor* j of i points its gradient link (given h_i) into i equals:

$$p_{j \rightarrow i} |_{h_i} = \left[\int_{h_i}^1 dh' \eta(h') \right]^{m_j} \equiv [\phi(h_i)]^{m_j}, \quad (1)$$

where $\phi(x) = \int_x^1 dh \eta(h)$ and m_j is the degree of node j . This is because the scalars at all the m_{j-1} neighbors of j (that do not include i) must be larger than h_i , and in addition we also have to have $h_j > h_i$. Hence, the probability that node j *does not* point its gradient link into i is $(1 - p_{j \rightarrow i} |_{h_i})$. In order for node i to be a leaf on the gradient tree, one must have that none of its neighbors point their gradient directions into it. This is given by: $\prod_{j=1}^{k_i} (1 - p_{j \rightarrow i} |_{h_i})$, where for simplicity we label the neighbors of i by $j = 1, 2, \dots, k_i$. Thus, the probability that i is a leaf is expressed by: $\pi_i = \int_0^1 dh_i \eta(h_i) \prod_{j=1}^{k_i} (1 - p_{j \rightarrow i} |_{h_i})$. Using (1) and the change of variable from h to ϕ , $d\phi = -dh \eta(h)$, the integral becomes:

$$\pi_i = \pi_i(m_1, m_2, \dots, m_{k_i}) = \int_0^1 d\phi \prod_{j=1}^{k_i} (1 - \phi^{m_j}). \quad (2)$$

This expression shows that among all nodes i with the same degree k , those that have neighbors with high degree are more likely to become leaves in $\nabla_h G$. It also shows that hubs will have very low probability of becoming leaves since in that case many (k) sub-unitary values are multiplied in equation (2). Therefore, leaves come from the set of nodes with low, or moderate degrees, connected to hubs, supporting figure 3(c). When there is relaxation, the leaves send their load towards higher degree nodes (on average), correlating the scalars with the degree. This increases the chance of larger degree nodes becoming leaves and breaking the star-like structures. As a consequence, the gradient network makes a transition to more elongated clusters, reducing its perimeter, and thus the congestion pressure. To support our claims numerically, we computed the average in-degree $\langle \ell(k) \rangle$ of $\nabla_h G$ as function of k in both the S and the R models. In the S model the gradient clusters are star-like, and thus they are very heterogeneous, which appears as a strong dependence of $\langle \ell(k) \rangle$ on k , see the circles in figure 4(a). After relaxation, however, this heterogeneity almost disappears (squares and triangles in figure 4(a)), and in the steady state, $\langle \ell(k) \rangle \approx 1$ (only a weak k -dependence remains), which means that $\nabla_h G$ is made of elongated, homogeneous trees.

Another way to understand the correlations between the scalar field and the local topology (degree) in the steady state of the dynamic model is by evaluating $h(k) - \langle h \rangle$ as function of the degree k , which is shown in figure 4(b) for different values of λ . It can be seen that as the degree increases the height increases as well. This means that in the steady state, higher degree nodes have higher scalar values, thus they are more likely to become leaves in the gradient network, a conclusion also supported by the Monte-Carlo studies of Gulbahce [18].




Figure 4. (a) Linear–log plot of $\langle \ell(k) \rangle$ as function of k with $\lambda = 2.5$ and $N = 1024$ for S (\circ), and R (\square), for $t = 0.1$ and R (\triangle) in the steady state. As the system evolves, $\langle \ell(k) \rangle$ loses its dependence on k . Simulations were averaged over 10^4 configurations. We show the plot in the linear–log scale in order to be able to visualize them in the same figure. (b) Linear–linear plot of $h(k) - \langle h \rangle$ as function of k for $N = 1024$ and different values of λ , $\lambda = 2.5$ (\circ), $\lambda = 3.0$ (\square) and $\lambda = 3.5$ (\triangle). The dashed line is used as a guide to show that for low k values, $h(k)$ is below $\langle h \rangle$. The simulations were done over 10^5 configurations.

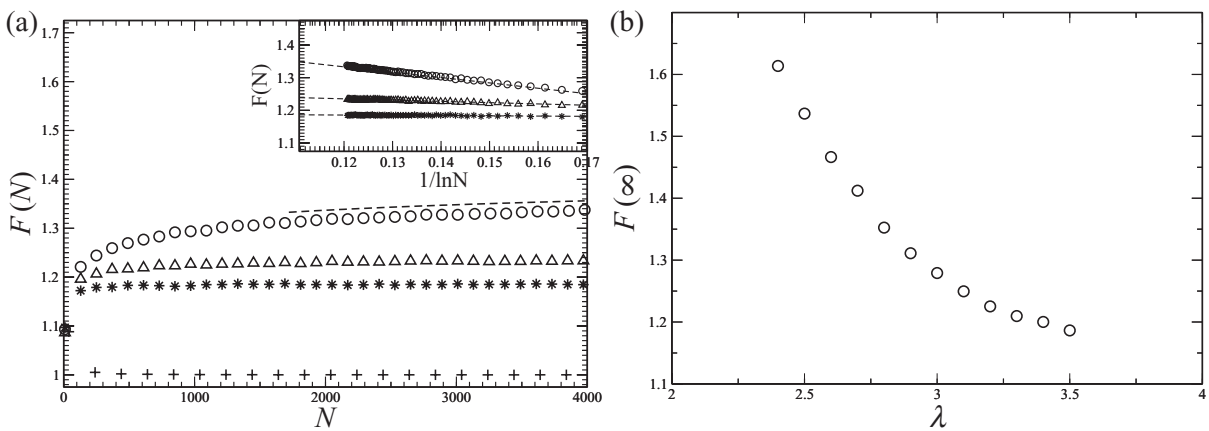


Figure 5. (a) $F(N)$ as function of N , for SF networks with $\lambda = 2.5$ (\circ), 3.0 (\triangle), 3.5 ($*$) and for ER networks with $p = 0.1$ ($+$). The dashed lines correspond to the logarithmic fit of $F(N)$ as function of N in the asymptotic regime (see also inset). (b) Asymptotic improvement factor as function of the connectivity exponent. Simulations were averaged over 10^3 network configurations.

Defining the relative improvement factor as $F(N) = J_S/J_R$, which measures how effective the relaxation dynamics is in reducing the congestion¹⁰, we find that for SF networks, as λ decreases, $F(N)$ increases, see figure 5(a), where we plot $F(N)$ as a function of N for different values of λ . In addition, from figure 5(a), we observe that $F(N)$ has a logarithmic convergence

¹⁰ Notice that $F > 1$ means that the congestion is reduced.

to its asymptotic value, i.e. $F(N) \simeq F(\infty) - K/\ln N$ (with K as a constant), see also the inset of figure 5(a). Figure 5(b) shows the infinite-system size relative improvement factor as function of λ , showing the increasing power of the relaxation mechanism for lower λ values. The effects shown here seem to hold for other SF networks that we tested as well, including real-world networks such as the DIMES Internet mapping project generated at the autonomous system (AS) level network (www.netdimes.org) [25]. The DIMES network is a correlated SF graph, and there the relaxation dynamics improves the congestion pressure drastically: from a $J_S = 0.61$ the congestion pressure drops to $J_R = 0.18$! This suggests that topology correlations play an important role which will be the subject of future studies. We observed similar behavior of the improvement factor for other real-world SF network topologies, as well.

In summary, we have shown that by introducing a simple and natural ‘surface’ relaxation mechanism, congestion in SF networks can considerably be reduced. Interestingly, the same mechanism has hardly any effect on congestion in the case of ER random graphs. Within the class of SF networks, the relaxational mechanism is more effective for lower λ values (however, still with $\lambda > 2$). These findings can be interpreted by tracking the morphological changes in the clusters of the corresponding gradient networks. Our results seem to hold for other SF networks as well, including real-world networks.

Acknowledgments

This work was supported by UNMdP and FONCyT (PICT 2005/32353). ZT, LAB and EL were supported in part by the US DOE grant W-7405-ENG-36.

References

- [1] Albert R and Barabási A-L 2002 *Rev. Mod. Phys.* **74** 47
Dorogovsetv S N and Mendes J F F 2003 *Evolution of Networks* (Oxford: Oxford University Press)
Caldarelli G 2007 *Scale-Free Networks* (Oxford: Oxford University Press)
- [2] Pastor-Satorras R and Vespignani A 2004 *Structure and Evolution of the Internet* (Cambridge: Cambridge University Press)
- [3] Barrat A, Barthelemy M and Vespignani A 2004 *Phys. Rev. Lett.* **92** 228701
- [4] Li L, Alderson D L, Doyle J and Willinger W 2006 *Internet Math.* **2** 431
Alderson D L and Li L 2007 *Phys. Rev. E* **75** 046102
- [5] Toroczkai Z and Bassler K E 2004 *Nature* **428** 716
- [6] Ohira T and Sawatari R 1998 *Phys. Rev. E* **58** 193
- [7] Fuks H and Lawniczak A T 1999 *Math. Comput. Simul.* **51** 101
- [8] Arenas A *et al* 2001 *Phys. Rev. Lett.* **86** 3196
- [9] Guimera R *et al* 2002 *Phys. Rev. Lett.* **89** 248701
- [10] Toroczkai Z *et al* 2008 *J. Phys. A: Math. Theor.* **41** 155103
- [11] Echenique P *et al* 2004 *Phys. Rev. E* **70** 056105
- [12] Danila B *et al* 2006 *Phys. Rev. E* **74** 046114
- [13] Sreenivasan S *et al* 2007 *Phys. Rev. E* **75** 036105
- [14] Rabani Y, Sinclair A and Wanka R 1998 *Proc. 39th Symp. on Foundations of Computer Science (FOCS)* pp 694–703
- [15] Valverde S and Solé R V 2004 *Eur. Phys. J. B* **38** 245
- [16] Korniss G *et al* 2003 *Science* **299** 677
- [17] Bollobás B 2001 *Random Graphs* 2nd edn (Cambridge: Cambridge University Press)

- [18] Gulbahce N 2007 *Chaos* **17** 026105
- [19] Family F 1986 *J. Phys. A: Math. Gen.* **19** L441
- [20] Pastore y Piontti A L, Macri P A and Braunstein L A 2007 *Phys. Rev. E* **76** 046117
- [21] Molloy M and Reed B 1995 *Comb. Prob. Comput.* **6** 161
- [22] Toroczkai Z *et al* 2006 *Computational Complexity and Statistical Physics* ed A Percus, G Istrate and C Moore (Oxford: Oxford University Press) p 249
- [23] Korniss G 2007 *Phys. Rev. E* **75** 051121
- [24] La Rocca C E, Braunstein L A and Macri P A 2008 *Phys. Rev. E* **77** 046120
- [25] Shavitt Y and Shir E 2005 DIMES—letting the Internet measure itself *Comput. Commun. Rev.* **35** 71–4