



This is the accepted manuscript made available via CHORUS. The article has been published as:

# Price of anarchy is maximized at the percolation threshold

Brian Skinner

Phys. Rev. E **91**, 052126 — Published 15 May 2015

DOI: [10.1103/PhysRevE.91.052126](https://doi.org/10.1103/PhysRevE.91.052126)

# The price of anarchy is maximized at the percolation threshold

Brian Skinner

*Materials Science Division, Argonne National Laboratory, Argonne, IL 60439, USA*

(Dated: February 15, 2015)

When many independent users try to route traffic through a network, the flow can easily become suboptimal as a consequence of congestion of the most efficient paths. The degree of this suboptimality is quantified by the so-called “price of anarchy” (POA), but so far there are no general rules for when to expect a large POA in a random network. Here I address this question by introducing a simple model of flow through a network with randomly-placed “congestible” and “incongestible” links. I show that the POA is maximized precisely when the fraction of congestible links matches the percolation threshold of the lattice. Both the POA and the total cost demonstrate critical scaling near the percolation threshold.

## I. INTRODUCTION

Optimizing the flow through a large network is no easy task. This fact is readily apparent to anyone who has lost their internet connection or been stuck in a traffic jam. The problem in such cases, and in a myriad of others, is that network links become congested when a large demand is placed on them, and this congestion lowers the network efficiency for all users. It is thus common for situations to arise where the self-interest of individual users, who would prefer to route their traffic through the most convenient links, is not aligned with the globally optimum behavior for all users, which generally requires that traffic be distributed more evenly to avoid congestion. Understanding this self-interest-driven inefficiency is critical for a large number of problems across an array of fields, including the planning of computer networks[1], transportation networks[2–5], and power grids[6], the allocation of public services[7], and even sports strategy[8, 9].

Unfortunately, at present there are no general rules for predicting which conditions give rise to large inefficiency in a random network. Indeed, previous works have largely studied congestible networks either at the level of proving general theorems[10–12] or by trying to model specific real-world systems as accurately as possible[4, 13]. In this paper, I seek a middle ground approach by introducing and analyzing a simple lattice model of randomly-placed congestible and incongestible links. The model takes inspiration from the spectacular history of similarly idealized lattice models in theoretical physics, which have provided important and generic insights for fields as diverse as electrical conduction in disordered materials[14–16], the spread of epidemics[17, 18], the mechanics of biopolymer networks[19], and the quantum Hall effect[20]. Crucially, in each of these examples the lattice percolation threshold plays a key role in determining the behavior of the system. As I show below, the same is true for the model proposed here.

For the sake of concreteness, the remainder of this paper discusses the network using the language of cars on a road network[21], but one can think of the words “roads,” “traffic,” and “commute time,” as equivalent to the more

generic concepts of “links,” “current,” and “cost.”

## II. PROPOSED MODEL

The canonical example of network suboptimality arising from user self-interest was introduced by the economist Arthur Pigou[22], who in 1920 imagined the simple network of two parallel roads shown in Fig. 1(a). In Pigou’s example, one of the two roads offers a constant commute time  $c_1 = 1$  (in some units). The second road offers a potentially much faster commute but is highly congestible. This road is described by a usage-dependent commute time  $c_2 = x_2$  (in the same units), where  $x_2 \in [0, 1]$  denotes the proportion of the total traffic on this second road. It is straightforward to show that the average commute time per driver,  $C = x_1 c_1 + x_2 c_2(x_2) = x_1 + x_2^2$ , where  $x_1 = 1 - x_2$  is the proportion of the traffic taking the first road, is optimized when the two roads are used equally:  $x_1 = x_2 = 1/2$ . This optimum arrangement has an average commute time  $C_{\text{opt}} = 3/4$ .

The optimum configuration, however, is unstable in a social sense: any driver on the first road experiences a commute time of  $c_1 = 1$ , while drivers on the second road experience a twice-shorter commute  $c_2 = 1/2$ . Thus, in the absence of any regulation, all drivers on the first road have an incentive to switch to the second road, and the resulting “equilibrium” has  $x_1 = 0$ ,  $x_2 = 1$ , and  $C = C_{\text{eq}} = 1$ . This equilibrium is equivalent to the Nash equilibrium from game theory, which is generally defined as the state in which no user can improve his commute time by changing his commute path [3].

The inefficiency of the equilibrium is quantified by the so-called “price of anarchy” (POA),  $P = C_{\text{eq}}/C_{\text{opt}} \geq 1$ ; one can think that the POA reflects the increase in global cost that results from allowing users to choose their own paths rather than having the optimum arrangement dictated to them.

While Pigou’s example provides an illustration of self-interest-driven inefficiency, there is no general rule for when to expect a significant POA in a large network. Motivated by this lack of qualitative understanding, I introduce a generalization of Pigou’s example in which




FIG. 1. (Color online) (a) Pigou's example of flow through a network with two "roads," one of which has constant commute time and the other of which is congestible. (b) The model introduced here, where constant roads (thin red lines) and congestible roads (thick blue lines) are randomly mixed. A unit amount of traffic is passed through the network, and all roads are unidirectional and proceed to the right.

slow "constant roads," having  $c(x) = 1$ , are combined with faster "congestible roads," having  $c(x) = x$ , to form a large network. The majority of this paper focuses on the case where roads are arranged into a two-dimensional square lattice, as shown in Fig. 1(b). All roads in the lattice are taken to be uni-directional, and a unit amount of traffic is passed through the network, such that all traffic enters through the  $L$  vertices on the left boundary and exits through the  $L$  vertices on the right boundary. The lattice is taken to be periodic in the vertical direction, and in this configuration each of the  $2^{2L}$  possible paths through the network involves traversing exactly  $2L$  roads. Each road type is assigned randomly, with a probability  $p$  of being of the congestible type.

The main focus of this paper is on calculating the global average commute time,

$$C = \sum_{\text{all roads } i} x_i c_i(x_i), \quad (1)$$

for both the optimum and equilibrium configurations as a function of  $p$  and the system size  $L$ . As outlined above, the traffic  $x_i$  on a given road  $i$  satisfies  $0 \leq x_i \leq 1$  and the commute time across the road is described by either  $c_i(x_i) = 1$  or  $c_i(x_i) = x_i$ .

A nontrivial dependence of the POA as a function of  $p$  can be anticipated by looking at the extreme cases of  $p = 0$  and  $p = 1$ . At  $p = 0$ , all roads have constant cost  $c = 1$  regardless of the arrangement of traffic, and so  $C_{\text{eq}} = C_{\text{opt}} = 2L$  and  $P(0) = 1$ . On the other hand, at  $p = 1$  all roads have  $c = x$ . In this case the symmetry of the lattice demands that all roads have equal usage  $x = 1/2L$  in both the equilibrium and optimum conditions, and as a consequence  $C_{\text{eq}} = C_{\text{opt}} = 1$  and  $P(1) = 1$ . For the case  $0 < p < 1$  however, where the network is non-uniform, one can expect  $P - 1$  to be finite, as fast

but congestible roads get overused in the equilibrium.

### III. NUMERICAL MODELING AND ELECTRICAL CIRCUIT ANALOGY

In order to find the total average commute time  $C$  in either the equilibrium or optimum situation, one should first solve for the traffic  $x_i$  along each road  $i$ . For the optimum case, this solution amounts to finding the minimum of the quadratic function  $C$  with respect to all variables  $x_i$  [see Eq. (1)], subject to the constraints that  $x_i \geq 0$  for all  $i$ , that the current is conserved at each node, and that the total current passed through the network is equal to unity. Similarly, the equilibrium can be thought of as the minimum of the modified cost function  $\sum_i \int_0^{x_i} c_i(x') dx'$ , subject to the same constraints. This equilibrium is equivalent to the state where all possible commute paths across the network have either the same commute time or have zero usage. Since finding the equilibrium or optimum therefore represents a minimization of a convex second-degree polynomial over a convex domain, the solution can be found using standard algorithms from convex optimization.[23] Details about the algorithm used here are given below.

Some insight into the problem of solving for the network flow can be gained by recasting it in the more familiar language of a linear electrical circuit. In this analogy, the traffic  $x$  along a given road is equated with a "current", and the commute time  $c(x)$  is equated with a "voltage drop". Thus, the constant roads with  $c(x) = 1$  are equivalent to links with unit voltage sources, which produce a current-independent voltage drop across the link. The congestible roads with  $c(x) = x$  are analogous to links with unit resistors, which produce a voltage drop that increases linearly with the current. Importantly, each link also includes an (ideal) diode [as depicted in Fig. 2(c)], which enforces the unidirectionality of the current. Such diodes ensure that the voltage drop across each link is strictly positive, or, equivalently, that all roads require a positive amount of time to traverse.

Once this analogy is made, the current on each link can be solved for using the corresponding Kirchoff laws. In particular, the Kirchoff loop law states that for any two paths  $A$  and  $B$  connecting the same two points in the network, the total voltage drops  $c_A$  and  $c_B$  along the paths are equal. This is equivalent to the Nash equilibrium condition, which states that any two paths with finite usage must have the same commute time; otherwise commuters will switch from the slower path to the faster one.

On the other hand, the optimum currents  $x_A$  and  $x_B$  along the two paths are not those which equalize  $c_A$  and  $c_B$ , but are instead those which minimize the average commute time,  $x_A c_A + x_B c_B$ , subject to the constraint that  $x_A + x_B$  is fixed. For nontrivial solutions having  $x_A, x_B > 0$ , this corresponds to the condition  $\partial_{x_A}(x_A c_A) = \partial_{x_B}(x_B c_B)$ . Consider, for example, that

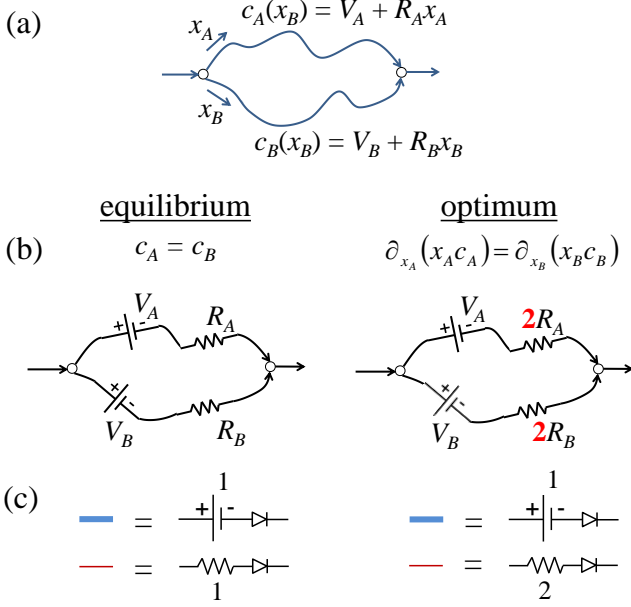


FIG. 2. (Color online) Electrical circuit analog of the lattice shown in Fig. 1(b). (a) Traffic splits between two paths, which generically have commute times  $c_A$  and  $c_B$ , respectively, that are linear functions of the traffic along each road. (b) The equilibrium traffic, for which  $c_A = c_B$ , is analogous to the current through a resistor network with resistances matching the value of  $\partial c / \partial x$  along each branch. The optimum traffic, on the other hand, is analogous to the current resulting from a circuit with *doubled* values of the resistance. Currents are assumed to proceed to the right. (c) The network model being considered [as in Fig. 1(b)] can generically be mapped onto an electrical circuit by representing constant roads as unit voltage sources and congestible roads as either unit resistors, for the equilibrium case, or as resistors with resistance 2, for the optimum case. The diodes ensure that current proceeds to the right, and that all commute times are positive.

the commute time along path A is described by the linear function  $c_A(x_A) = V_A + R_A x_A$ , where  $V_A$  and  $R_A$  are coefficients that represent the number of constant or congestible roads, respectively, along path A. Let the commute time along path B be similarly written  $c_B(x_B) = V_B + R_B x_B$ . In this case it is easy to show that a nontrivial optimum satisfies  $V_A + 2R_A x_A = V_B + 2R_B x_B$ . But this relation is precisely the Kirchhoff loop law for the case where the “resistances”  $R_{A/B}$  are doubled. Thus, one can solve for the optimum currents simply by modeling each congestible road as a resistor with resistance 2 rather than as a unit resistor, as depicted in Fig. 2(b). [Note, for example, that such a substitution gives the optimum currents  $x_A = x_B = 1/2$  in Pigou’s example, Fig. 1(a).] This result, while elementary, already has important implications for network routing: it implies that traffic through a linear, congestible network is optimized when users respond to a cost that rises twice as fast with increased usage as does the bare transit time cost. The possibility of achieving such a cost rate doubling in real

networks is already being studied; for example, on major roadways it might be achievable through through dynamic tolling.[13]

Of course, for the problem considered here, the process of solving for the circuit currents is complicated by the presence of diodes on each link (or, equivalently, by the constraint that all commute times be strictly positive). This condition necessitates the use of a more careful procedure than a simple solution of Kirchhoff’s equations, since solving for the current in all circuit elements requires one to simultaneously ascertain the correct state of all  $(2L)^2$  diodes. To find this solution, I use here a simple adaptation of the “greedy” algorithm that is commonly used to search for the ground state of spin or Coulomb glasses[24].

In particular, in this algorithm the initial state of each diode (“on” or “off”) is first guessed randomly, and the corresponding system of Kirchhoff equations is solved numerically. The resulting currents and voltages coming from this initial solution have, in general, a number of violations of the assumed states of the diodes. That is, some diodes assumed to be “on” have negative currents in the (erroneous) solution, and some diodes assumed to be “off” have positive voltage drops across them. In the greedy algorithm, the “on” diode with the largest negative current and the “off” diode with the largest positive voltage both have their states switched, and the Kirchhoff laws are re-solved. This process is continued until a solution is reached that has complete consistency between the currents and the diode states. A uniqueness theorem for linear circuits with diodes[25] guarantees that any such solution is unique in terms of its dissipated power, which is the analogue of the total commute time  $C$ . If no such solution is reached after a large number of numerical iterations (which happens rarely), then the process is re-initialized using a different initial guess for the diode states.

This procedure is implemented numerically for each random realization of the network in order to calculate values of the traffic  $\{x_i\}$  on all of the system’s roads in both the optimum and equilibrium situations. The resulting total commute times,  $C_{\text{opt}}$  and  $C_{\text{eq}}$ , are then calculated using Eq. (1). Results presented below correspond to averages over many random networks for each value of  $p$  and each system size  $L$ .

#### IV. RESULTS

The results for the POA are shown in Fig. 3 as a function of  $p$  for a range of system sizes. Notably, the POA acquires its maximum value very close to the point where  $p$  matches the percolation threshold for directed percolation through the lattice,  $p_c \approx 0.6447$ . [26] As the system size is increased, the maximum of  $P$  becomes increasingly sharp and moves closer to  $p = p_c$ . As shown below, scaling of the curves  $P(p, L)$  suggests that in the limit of infinite system size, the POA peak is infinitely sharp and

is located precisely at  $p = p_c$ .




FIG. 3. (Color online) POA as a function of the fraction  $p$  of congestible roads in the lattice, plotted for different values of the system size  $L$ . The vertical dotted line indicates the lattice percolation threshold,  $p_c \approx 0.6447$ .

While the focus of this paper is on the square lattice shown in Fig. 1(c), I briefly note that one can verify the generality of the main result of this paper, that the POA is maximized at the percolation threshold, by examining other lattice types with different values of  $p_c$ . Such lattices can be expected to produce qualitatively similar curves  $P(p)$  as in Fig. 3, but with the maximum shifted to the percolation threshold of the lattice being considered. This is demonstrated explicitly in Appendix A for one specific case.

That the POA achieves a maximum at the percolation threshold can be rationalized using the following qualitative argument. At  $p < p_c$ , there are no continuous paths in a large system that connect opposite faces of the system while traversing only congestible roads, as illustrated in Fig. 4(a). Thus, all traffic across the network must use a combination of congestible roads and the slower constant roads. In this situation the equilibrium traffic is naturally distributed over many paths with relatively low susceptibility to congestion, and the POA is not too large. Exactly at the percolation threshold,  $p = p_c$ , there appears a single macroscopic pathway (the “infinite cluster,” in percolation language[24]) connecting opposite sides of the system that uses only congestible roads. This single pathway becomes heavily over-used in the equilibrium, and the POA is relatively large.

Finally, when  $p > p_c$  there are many pathways connecting opposite faces that use only congestible roads, as shown in Fig. 4(b), and as  $p$  is increased the number of such pathways increases. At such large  $p$  the constant roads are abandoned in the equilibrium, and the POA is a reflection only of the degree of congestion on those pathways that have finite equilibrium usage. Increasing  $p$  eases that congestion, and the POA falls abruptly. Fig. 5 shows a visualization of the traffic through the network at  $p \approx p_c$  in both the equilibrium and optimum config-




FIG. 4. (Color online) Schematic illustration of large clusters and percolating pathways of congestible roads. (a) At  $p < p_c$  and  $|p - p_c| \ll 1$ , the system contains large, disconnected clusters of congestible roads, with typical size  $\xi_{\parallel}$  in the down-stream direction and  $\xi_{\perp}$  in the perpendicular direction. One such cluster is highlighted in orange. (b) At  $p$  slightly larger than  $p_c$ , on the other hand, there are many parallel pathways for traversing the lattice using only congestible roads. The correlation lengths  $\xi_{\parallel}$  and  $\xi_{\perp}$  describe the typical horizontal and vertical separation between these paths. Small, isolated clusters and “dead ends” are not shown.

urations. As one can see, the traffic is distributed over more paths in the optimum than in the equilibrium. Additional visualizations of the traffic density are given in Appendix B.




FIG. 5. (Color online) Spatial map of the traffic flow in a typical realization of a random network at  $p \approx p_c$ . Roads that carry finite traffic in the equilibrium are shown as black points, and light blue points indicate additional roads that are used in the optimum. The image has been cropped vertically, but shows the full width ( $L = 75$ ) of the system.

In order to understand the scaling behavior of the equilibrium and optimum commute times, let us first consider the case where the system size  $L \rightarrow \infty$ , while  $p - p_c$  remains finite. At  $p < p_c$  and  $|p - p_c| \ll 1$ , the network contains large but disconnected clusters of fast, congestible roads with maximum size  $\xi_{\parallel} \propto |p - p_c|^{-\nu_{\parallel}}$  in the down-



stream direction and  $\xi_\perp \propto |p - p_c|^{-\nu_\perp}$  in the perpendicular direction [27, 28], as illustrated in Fig. 4(a). Here,  $\nu_\parallel \approx 1.733$  and  $\nu_\perp \approx 1.097$  are critical exponents [26], with  $\nu_\parallel > \nu_\perp$  indicating larger correlation length in the downstream direction. (This behavior can also be seen qualitatively in Fig. 5, which shows that near the percolation threshold large, asymmetric holes appear in the current paths.)

As traffic passes from the left to the right side of the network, it generally seeks to avoid the slow constant roads, and therefore it preferentially follows the “backbone” of large percolation clusters. However, since these clusters are not connected macroscopically, the traffic must pass through at least one constant road each time it moves from one cluster to another. Therefore, the traffic generally passes through  $\sim 1$  such road for each path length  $\xi_\parallel$  travelled. Since only a unit amount of current is passed through the entire system, the amount of traffic through any given percolation cluster is vanishingly small in the limit  $L \rightarrow \infty$ , and the commute time across the congestible roads within the cluster is also vanishingly small. Consequently, the total commute time  $C$  is dominated by passage through the constant roads connecting adjacent clusters. Since commuting across the entire lattice requires drivers to pass from one large cluster to another  $\sim L/\xi_\parallel$  times, one can expect  $C$  to scale as

$$C \sim \frac{L}{\xi_\parallel} \sim L|p - p_c|^{\nu_\parallel} \quad (2)$$

at  $p < p_c$ .

On the other hand, at  $p > p_c$ , there are many parallel pathways for traversing the lattice that use only congestible roads, as depicted in Fig. 4(b). Each such pathway takes only a small fraction  $x' \sim \xi_\perp/L$  of the total current, while paths that traverse constant roads are completely abandoned. Since the total commute time along any given percolating pathway is proportional to the commute time  $c(x') = x'$  on a given congestible road multiplied by the total path length  $\sim L$ , one can say that the typical average time  $C \sim x'L$ . In other words,

$$C \sim \xi_\perp \sim |p - p_c|^{-\nu_\perp} \quad (3)$$

at  $p > p_c$ . Equations (2) and (3) are shown together with numerical results for  $C$  in Fig. 6.

One can notice that Eqs. (2) and (3) represent very different behavior: at  $p < p_c$  the commute time scales extensively with the system size,  $C \propto L^1$ , while at  $p > p_c$  the commute time becomes independent of system size,  $C \propto L^0$ . Exactly at the threshold,  $p = p_c$ , one can expect the commute time to scale as some nontrivial power of the system size:  $C \propto L^m$ . An estimate for this exponent  $m$  can be obtained by equating Eqs. (2) and (3) and solving for the corresponding value of  $|p - p_c|$  at which the two relations cross over to each other. This procedure gives

$$C(p = p_c) \sim L^{\nu_\perp/(\nu_\parallel + \nu_\perp)}, \quad (4)$$




FIG. 6. (Color online) Dependence of the commute time on the fraction  $p$  of congestible roads, plotted for  $L = 50$ . The solid (blue) line shows  $C_{\text{opt}}$ , and the (red) dots are  $C_{\text{eq}}$ . The dashed lines show the analytical results of Eqs. (2) and (3), respectively, for  $L \rightarrow \infty$ . The dotted vertical line indicates  $p = p_c$ .

or  $m = \nu_\perp/(\nu_\parallel + \nu_\perp) \approx 0.388$ . This analytical estimate for  $m$  is consistent with numerical results, as shown in Fig. 7. An independent fitting of the data in Fig. 7 gives  $m \approx 0.31$ . This same scaling relation  $C \propto L^m$  describes both the equilibrium and optimum. (Indeed, general theorems have shown that  $C_{\text{eq}}$  and  $C_{\text{opt}}$  cannot differ by more than a constant numeric factor [11]. For networks with linear cost functions,  $C_{\text{eq}}/C_{\text{opt}} \leq 4/3$ .)




FIG. 7. (Color online) Scaling of the average commute time  $C$  with the system size  $L$ . Different curves are labelled by their corresponding value of  $p$ . At  $p < p_c$ , the commute time scales as  $C \propto L^1$ , and at  $p > p_c$  the commute time scales as  $C \propto L^0$ . Precisely at  $p = p_c$ , the commute time follows  $C \propto L^m$ , with  $m \approx 0.388$ . Filled dots denote  $C_{\text{eq}}$ , and open circles are  $C_{\text{opt}}$ .

For finite values of the system size, and for  $p$  close to the percolation threshold, it should be possible to write the scaled commute time  $C/L^m$  using a critical scaling form. Such scaling usually takes the form  $F = f(L/\xi)$ ,

where  $F$  is some system property and  $\xi \propto |p - p_c|^{-\nu}$  is the correlation length, so that  $F$  is a function only of  $L/\xi$ , or equivalently of  $(L/\xi)^{1/\nu} = (p - p_c)L^{1/\nu}$ . Such scaling is indeed possible here, as demonstrated explicitly in Fig. 8(a). In this figure the scaled cost  $C/L^m$  is shown to be a function only of the combination  $(p - p_c)L^{1/\nu}$ , with  $\nu \approx 2.62 \pm 0.2$  determined by best fit of the data collapse[29]. Fig. 8(b) demonstrates that this same value of  $\nu$  produces good collapse of the curves from Fig. 3 for the POA. As mentioned above, such scaling suggests that for  $L \rightarrow \infty$  the POA peak converges precisely to  $p = p_c$ .

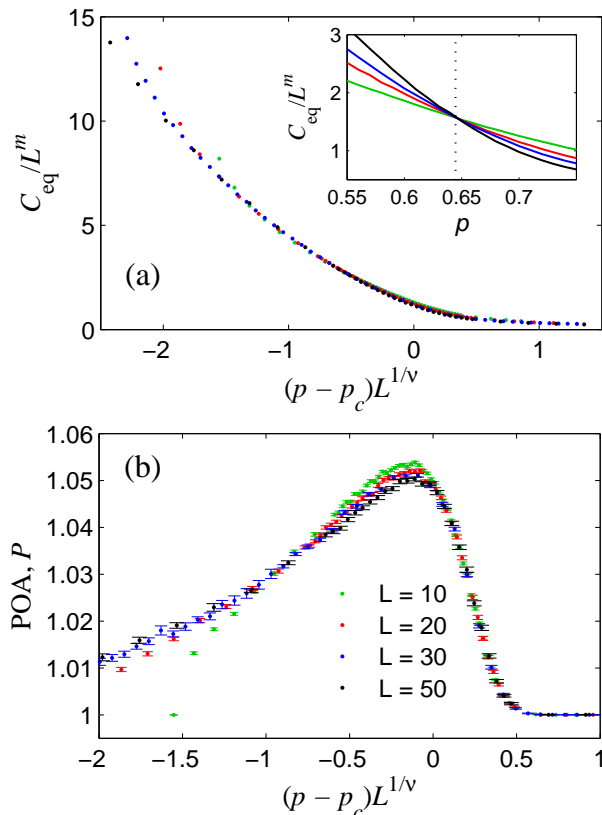


FIG. 8. (Color online) Critical scaling. (a) Inset: The scaled equilibrium commute time  $C_{\text{eq}}/L^m$  near the percolation threshold (vertical dotted line). Main figure:  $C_{\text{eq}}/L^m$  can be parameterized for all system sizes by  $(p - p_c)L^{1/\nu}$ , with  $\nu \approx 2.6$ .  $C_{\text{opt}}$  (not shown) can be similarly scaled. (b) Near the percolation threshold, different curves for the POA also collapse when plotted as a function of  $(p - p_c)L^{1/\nu}$ .

Of course, the problem being considered here has two separate correlation lengths, so that the fitted exponent  $\nu$  should in fact be a linear combination of the two critical exponents  $\nu_{\parallel}$  and  $\nu_{\perp}$ . One reasonable expectation is that the relevant scaling form is  $C/L^m = f(L/L_c)$ , where  $L_c \sim \xi_{\perp}\xi_{\parallel}$  is the system size for which Eqs. (2) and (3) become equal. In other words,  $L_c$  represents the system size below which the system cannot be unambiguously described as percolating or non-percolating, in terms of its average commute time. This hypothesis leads to the

conclusion that  $\nu = \nu_{\parallel} + \nu_{\perp} \approx 2.83$ , which is consistent with the result presented above. Future works can validate this hypothesis explicitly by studying systems with varying aspect ratio.

## V. CONCLUSION

In summary, this paper has introduced a simple model of random congestible networks and demonstrated a clear and previously unnoticed connection between percolation and self-interest-driven inefficiency. A number of generalizations and extensions of the model deserve further exploration, including an extension to other network topologies [10, 30], and to the case of nonlinear cost functions. For highly nonlinear cost functions one can, in general, expect a significantly larger value of the POA at the percolation threshold[31]. Cost functions that include a critical “jamming” density, as are realistic for actual highways [32], may give rise to a diverging POA.

More generally, this work hints at the possibility of a deeper connection between POA and percolation that may go well beyond simple Pigou-type models. Such a connection seems to provide a novel and interesting playground for statistical mechanics, and has the potential to provide invaluable understanding for efforts to mitigate network congestion effects across a wide set of disciplines.

## ACKNOWLEDGMENTS

I am grateful to A. Nahum, V. Sokolov, K. A. Matveev, T. Roughgarden, A. Glatz, and R. Brierley for helpful discussions, and to W. DeGottardi, G. Guzman-Verri, and A. Lopez-Bezanilla for their comments on the manuscript. Work at Argonne National Laboratory was supported by the U.S. Department of Energy, Office of Science under contract no. DE-AC02-06CH11357.

### Appendix A: POA in a different lattice type

As mentioned in the main text, one natural test of the generality of the claim that the POA is maximized at  $p = p_c$  is to analyze flow through a different type of lattice. Here I show results for the three-dimensional (3D) body-centered cubic (BCC) lattice, which can be thought of as the 3D generalization of the lattice in Fig. 1(b). The BCC lattice has a significantly smaller percolation threshold  $p_c \approx 0.2873$  due to its twice-larger coordination number. The POA as a function of  $p$  for this lattice is plotted in Fig. 9. As one can see, the behavior of  $P(p)$  is qualitatively similar to that of Fig. 3, but with the maximum shifted to the new percolation threshold.




FIG. 9. (Color online) POA for flow through a 3D BCC lattice with randomly-placed congestible and constant roads, plotted for different system sizes. The vertical dotted line shows the lattice percolation threshold,  $p_c \approx 0.2873$ .

### Appendix B: Traffic density maps

Figure 5 of the main text provides a visualization of the traffic through a particular realization of the network by showing which roads are used in the equilibrium and optimum situations. Here I provide an additional visualization of the traffic density.

In Fig. 10 are shown the equilibrium and optimum traffic densities  $x$  for each road of the same random network depicted in Fig. 5. In addition to having smaller “holes” in the traffic pattern (as illustrated in Fig. 5), the optimum also has a more even distribution of the traffic flow as compared to the equilibrium.






FIG. 10. (Color online) Traffic density in a network close to the percolation threshold,  $p \approx p_c$ . The color indicates the traffic on a particular road, normalized to the system-averaged traffic per road,  $x_{\text{avg}} = 1/(2L)$ . The equilibrium configuration (left) generally has its traffic concentrated on fewer roads, while the optimum (right) has a more even distribution.

- 
- [1] T. Roughgarden, *Selfish routing and the price of anarchy* (MIT press, 2005).
  - [2] D. Braess, *Zeitschrift für Operations-Research* **12**, 258 (1968).
  - [3] J. G. Wardrop, in *ICE Proceedings: Engineering Divisions*, Vol. 1 (Thomas Telford, 1952) pp. 325–362.
  - [4] H. Youn, M. T. Gastner, and H. Jeong, *Physical Review Letters* **101**, 128701 (2008).
  - [5] S. Zhu, D. Levinson, and H. Liu, in *Transportation Research Board 89th Annual Meeting Compendium of Papers*, Vol. 1005 (Citeseer, 2010) pp. 10–2298.
  - [6] D. Witthaut and M. Timme, *New Journal of Physics* **14**, 083036 (2012).
  - [7] V. A. Knight and P. R. Harper, *European Journal of Operational Research* **230**, 122 (2013).
  - [8] B. Skinner, *Journal of Quantitative Analysis in Sports* **6**, Article 3 (2010).
  - [9] B. Skinner and B. Carlin, *Significance* **10**, 25 (2013).
  - [10] G. Valiant and T. Roughgarden, *Random Structures & Algorithms* **37**, 495 (2010).
  - [11] T. Roughgarden and v. Tardos, *J. ACM* **49**, 236 (2002).
  - [12] T. Roughgarden, *Journal of Computer and System Sciences Special Issue on STOC 2002*, **67**, 341 (2003).
  - [13] T. Tsekeris and S. Vo, *NETNOMICS: Economic Research and Electronic Networking* **10**, 5 (2009).
  - [14] R. Fisch and A. B. Harris, *Physical Review B* **18**, 416 (1978).
  - [15] A. Skol and B. Shklovskii, *Sov. Phys. Semicond.-USSR* **8**, 1029 (1975).
  - [16] V. L. Nguen, B. Z. Spivak, and B. I. Shklovskii, *Zh. Eksp. Teor. Fiz* **89**, 1770 (1985), [*JETP* 62, 1021, (1985)].
  - [17] C. Moore and M. E. J. Newman, *Physical Review E* **61**, 5678 (2000).
  - [18] L. M. Sander, C. P. Warren, I. M. Sokolov, C. Simon, and J. Koopman, *Mathematical Biosciences* **180**, 293 (2002).
  - [19] C. P. Broedersz, X. Mao, T. C. Lubensky, and F. C. MacKintosh, *Nature Physics* **7**, 983 (2011).
  - [20] J. T. Chalker and P. D. Coddington, *Journal of Physics C: Solid State Physics* **21**, 2665 (1988).
  - [21] Such language may become particularly relevant with the further development of self-driving cars[33].
  - [22] A. C. Pigou, *The economics of welfare* (Transaction Publishers, 1924).
  - [23] S. Dasgupta, C. H. Papadimitriou, and U. Vazirani, *Algorithms*, 1st ed. (McGraw-Hill, Inc., New York, NY, USA, 2008).
  - [24] B. I. Shklovskii and A. L. Efros, *Electronic Properties of Doped Semiconductors* (Springer-Verlag, New York, 1984).
  - [25] W. L. Black and J. H. Saltzer, *Tech Engineering News* **XLI**, 18 (1960).
  - [26] I. Jensen, *Journal of Physics A: Mathematical and General* **32**, 5233 (1999).
  - [27] F. Family, *Journal of Physics A: Mathematical and General* **15**, L583 (1982).
  - [28] S. Redner, *Physical Review B* **25**, 3242 (1982).
  - [29] S. M. Bhattacharjee and F. Seno, *Journal of Physics A: Mathematical and General* **34**, 6375 (2001).
  - [30] S. Hata, H. Nakao, and A. S. Mikhailov, *Physical Review E* **89**, 020801 (2014).
  - [31] Generally, a network whose cost functions are monotonically increasing and involve a power law of degree no larger than  $\alpha$  have  $P \leq (\alpha + 1)^{1+1/\alpha} / [(\alpha + 1)^{1+1/\alpha} - \alpha]$  [12].
  - [32] A. D. May, *Traffic Flow Fundamentals* (Prentice-Hall, Englewood Cliffs, NJ, 1990).
  - [33] J. Markoff, *The New York Times* **10**, A1 (2010).