The Price of Anarchy in flow networks as a function of node properties

O. Smith[1], J. Crowe[2], R.D. O'Dea[1] and K.I. Hopcraft[1]

[1] *School of Mathematical Sciences, University of Nottingham, United Kingdom, NG7 2RD*
[2] *Faculty of Engineering, University of Nottingham, United Kingdom, NG7 2RD*

PACS 89.75.-k – Complex systems
PACS 89.75.Hc – Networks and genealogical trees
PACS 89.90.+n – Other topics in areas of applied and interdisciplinary physics

Abstract – Many real-world systems such as traffic and electrical flow are described as flows following paths of least resistance through networks, with researchers often focusing on promoting efficiency by optimising network topology. Here, we instead focus on the impact of network node properties on flow efficiency. We use the Price of Anarchy $\mathcal{P}$ to characterise the efficiency of least-resistance flows on a range of networks whose nodes have the property of being sources, sinks or passive conduits of the flow. The maximum value of $\mathcal{P}$ and the critical flow volume at which this occurs are determined as a function of the network's node property composition, and found to have a particular morphology that is invariant with network size and topology. Scaling relationships with network size are also obtained, and $\mathcal{P}$ is demonstrated to be a proxy for network redundancy. The results are interpreted for the operation of electrical micro-grids, which possess variable numbers of distributed generators and consumers. The highest inefficiencies in all networks are found to occur when the numbers of source and sink nodes are equal, a situation which may occur in micro-grids, while highest efficiencies are associated with networks containing a few large source nodes and many small sinks, corresponding to more traditional power grids.

Introduction. – Flows on networks, such as traffic taking routes of shortest travel time or electrical current taking paths of least resistance though a network of connections can waste resources because they follow a local rather than a system-wide optimisation of the flow. For example, drivers generally behave non-cooperatively when selecting shortest routes, leading to traffic congestion that could be avoided by the intervention of a central management with a global perspective [1]. When agents compete selfishly for resources or to minimise their effort, they eventually attain a Nash equilibrium [2, 3], whereby any change in their strategy fails to further lower their costs. The Price of Anarchy $\mathcal{P}$ [4] gauges the inefficiency caused by this lack of cooperation [5] and is defined as the ratio of the cost of the worst Nash equilibrium to that of the system's global optimum (GO). In this Letter, $\mathcal{P}$ is established as a computationally efficient measure of inefficiency and network redundancy for flows such as electricity. The dependence of $\mathcal{P}$ on the numbers of flow sources and sinks, and network structure, is also addressed, and found to possess properties that are invariant with regard to networks of different topology.

$\mathcal{P}$ has been studied in a variety of contexts, such as in network growth games [6], job scheduling [7], resource allocation in public services [8], supply chains [9], and in network traffic flows where a cost (*i.e.* travel time) is incurred for traversing edges [10, 11]. If the individual drivers comprise only a very small amount of the overall flow, then it can be treated as a continuous quantity. Such flows also serve as a model for electrical current, comprising infinitesimally small particles, following paths of least resistance [12]. The Nash equilibrium corresponds to all routes on the network between an arbitrarily chosen source-sink pair having equal cost [13], or local voltage drop in the case of an electrical network, such that no change in flow pattern or routing can lower the cost. In [14] the upper bound on $\mathcal{P}$ was found to be 4/3 if the edge cost functions are linear functions of flow volume. Although these worst case values of $\mathcal{P}$ are independent of network topology, depending only on the class of edge function, values of $\mathcal{P}$ that differ from these extremes are strongly influenced by topology, flow volume, placement of sources and sinks and distribution of parameters in cost functions [4, 11]. For example, [15, 16] considered the case of a lattice network and revealed how $\mathcal{P}$ is affected by the size, aspect ratio and total flow through the lattice.

In some cases, the addition of new edges into a network can cause a counter-intuitive increase in the cost of the flow due to the inefficiency of the Nash equilibrium. This is referred to as Braess's paradox [17, 18], and has been studied in traffic networks [11, 19], where the addition of a road can increase average travel time, and electrical circuits [12]. Variants of this phenomenon have also been reported in supply chains [20] and oscillator networks [21, 22]; refer to [23] for an overview.

Previous studies of the Price of Anarchy have considered sources and sinks of flow only in specific arrangements. In [14] a single source-sink pair was considered whereas [11] treated ordered source-sink pairs with characteristic flows along overlapping paths occurring between them. The present Letter first establishes a connection between $\mathcal{P}$ and the efficiency and redundancy of least-resistance network flows, and then investigates the dependence of $\mathcal{P}$ on the relative and absolute numbers of flow source and sink nodes, to ascertain whether, for a given network, the configuration of node types can be altered to change efficiency. This is of importance to the design and control of electrical micro-grids which typically have varying numbers of low output intermittent sources of electrical power distributed throughout their structure. As the drive towards smaller, distributed generators becomes more urgent in order to mitigate climate change, understanding the impact of variable generation on electrical networks presents a pressing interdisciplinary challenge [24].

Network flow model. – We consider flows though graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ edges, wherein n_s node have the property of being sources of flow, n_d are sinks and the remaining n_p are passive or empty. Each edge $e \in \mathcal{E}$ has a linear cost function $c_e(f_e) = \alpha_e f_e + \beta_e$, where f_e is the volume of flow or electrical current on that edge. The functions c_e can be interpreted as the voltage drop across the edge, while the coefficients α_e and β_e represent Ohmic resistance and flow independent voltage drops respectively. For a flow vector $f \in \mathbb{R}^m$ the total cost across the network is $\mathcal{C}(f) = \sum_e c_e(f_e) f_e$, representing total power loss. The global optimum flow f_{GO} is then the flow pattern that minimises this cost:

$$\min_f \ \mathcal{C}(f) \quad \text{constrained by } Ef = b, \tag{1}$$

where $E \in \mathbb{R}^{n \times m}$ is the node-edge incidence matrix and b is the flow injection vector with components

$$b_v = \begin{cases} (1+\xi_v)F/n_s, & \text{if node } v \text{ is a source,} \\ -(1+\xi_v)F/n_d, & \text{if node } v \text{ is a sink,} \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

with F being the total flow or current injected into the network, and ξ_v being random noise. The condition $Ef = b$ enforces conservation of flow at nodes, equivalent to Kirchoff's current law. The Nash equilibrium flow f_{Nash} is given by the optimisation problem [13]

$$\min_f \sum_e \int_0^{f_e} c_e(q)\,\mathrm{d}q \quad \text{constrained by } Ef = b. \tag{3}$$

The optimisation problems in (1) and (3) are both convex and solved using subgradient projection methods [26]. The Price of Anarchy is then $\mathcal{P} = \mathcal{C}(f_{\mathrm{Nash}})/\mathcal{C}(f_{\mathrm{GO}}) \equiv \mathcal{C}_{\mathrm{Nash}}/\mathcal{C}_{\mathrm{GO}}$.

Nash equilibria conditions are equivalent to Kirchoff's voltage law. – A physical interpretation of the Nash equilibria obtains from a consideration of Kirchoff's voltage law (KVL), which states that voltages around closed cycles in an electrical network sum to zero. If there is a cycle embedded in a network, then there will be at least two distinct paths between a pair of source and sink nodes. At the Nash equilibrium, each arm of the cycle must have equal cost; hence the cost of any traversal around the cycle is zero, and so the Nash equilibrium condition is equivalent to KVL. The Nash flow therefore necessarily satisfies both Kirchoff's current and voltage laws and is thus a physically legitimate electrical flow for an electrical network in stable operation with matched supply and demand. The relative inefficiency of this flow, resulting in $\mathcal{P} > 1$, stems from the constraints of Kirchoff's conservation laws that define the Nash equilibrium.

Relationship with network redundancy. – $\mathcal{P}$ measures the disparity between the costs associated with the Nash and GO flows. In an electrical context the GO would correspond to a flow being able to violate KVL in order to minimise total power loss; however, such an equilibrium would nevertheless be desirable to obtain because it minimises the power consumed by the network. Therefore, $\mathcal{P}$ remains a useful metric for assessing efficiency in networks with flows following paths of least resistance, and also for topological redundancy as we now show.

Consider the network shown in fig. 1(a), first introduced by Pigou [27], being the smallest graph admitting a value of $\mathcal{P} > 1$, and which serves as the canonical example to demonstrate the Price of Anarchy [11, 13]. Edge 1 has variable cost $c_1 = f_1$, whereas edge 2 has fixed cost $c_2 = 1$. F units of flow enter on the left and exit on the right. Fig. 1(b) shows the value of $\mathcal{P}$ in this network as a function of F. For $0 < F \leq 1/2$, indicated by the unshaded area, all flow is routed over edge 1 under both the Nash and GO equilbria, with identical costs $\mathcal{C} = F^2$; consequently $\mathcal{P} = 1$. For $1/2 < F \leq 1$ (light gray area), $f_1 = F$ under the Nash flow, so $\mathcal{C}_{\mathrm{Nash}} = F^2$. The GO minimises its cost when $f_1 = 1/2$, $f_2 = F - 1/2$ and the total cost is then $\mathcal{C}_{\mathrm{GO}} = F - 1/4$, giving $\mathcal{P} = F^2/(F - 1/4)$. For $F > 1$ (dark gray area), the Nash equilibrium routes all flow surplus of 1 through edge 2, giving $\mathcal{C}_{\mathrm{Nash}} = F$, whereas the GO remains unchanged – hence, $\mathcal{P} = F/(F - 1/4)$.

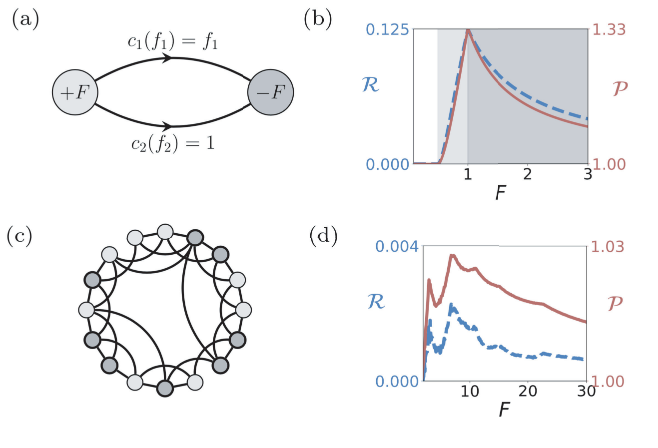

Fig. 1: (a) The example Pigou network comprising a source and sink node connected by a variable and fixed cost edge. (b) $\mathcal{P}$ (red line) and $\mathcal{R}$ (blue dashed line) for the Pigou network in (a) are shown as functions of F. (c) A small world network with $q = 0.1$, $k = 4$, $n = 16$, $n_s = n_d = 8$, $n_p = 0$. (d) $\mathcal{P}$ and $\mathcal{R}$ shown as functions of F for the small world network shown in (c).

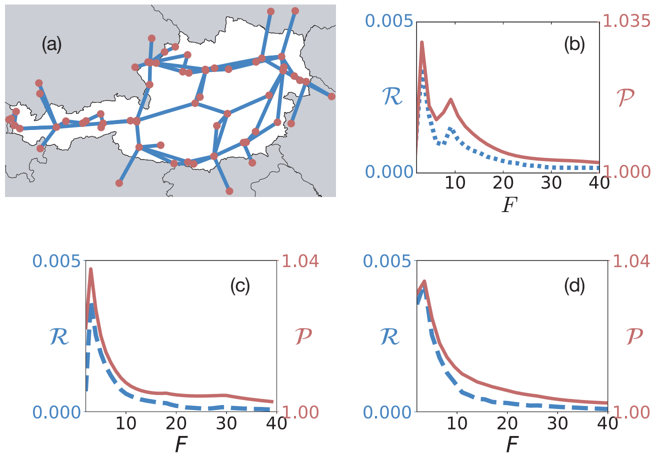

Fig. 2: (a) The Austrian power grid, constructed from open source topological data from [29]. (b) $\mathcal{P}$ and $\mathcal{R}$ for the network in (a) as a function of total current F, where F has been normalised using the per-unit system. (c) and (d) show $\mathcal{P}$ and $\mathcal{R}$ for the IEEE 14 bus and 118 bus test networks respectively, where the flow F has again been normalised into the per-unit system,

We now establish a qualitative relationship between $\mathcal{P}$ and network redundancy. Recall that the Nash equilibrium condition and KVL are equivalent in electrical networks. It is possible to drive the Nash flow, with cost $\mathcal{C}_{\text{Nash}}$, towards the GO by manipulating the network such that excess flow is transferred from edge 1 to edge 2. This is achieved by reducing the capacity on edge 1. This excess capacity is given by the difference between the flow on edge 1 for each equilibrium, i.e. $f_1^{\text{Nash}} - f_1^{\text{GO}}$. The equilibrium on this modified network has cost $\mathcal{C}'_{\text{Nash}} \leq \mathcal{C}_{\text{Nash}}$. This means that edge 1 provides redundant capacity that can be removed. Defining this edge redundancy in terms of the costs obtains:

$$\mathcal{R}^e = \frac{\mathcal{C}_{\text{Nash}} - \mathcal{C}'_{\text{Nash}}}{\mathcal{C}_{\text{Nash}}} = \begin{cases} 0, & 0 \leq F < 1/2, \\ (F - 1/2)^2 / F^2, & 1/2 \leq F < 1, \\ 1/4F, & F > 1, \end{cases} \tag{4}$$

which is the relative decrease in cost available by removing capacity from edge 1. No relative decrease in cost is possible by removing any capacity from edge 2. In order to generalise this measure to larger networks it is averaged over both edges to give $\mathcal{R} := \overline{\mathcal{R}^e}$, which is the mean decrease in cost attainable by removing capacity from an edge. Fig. 1(b) shows $\mathcal{R}$, whose form emulates $\mathcal{P}$. For larger and more complex networks, such as the small world network depicted in fig. 1(c), this correspondence between $\mathcal{P}$ and $\mathcal{R}$ prevails, as shown in fig. 1(d).

The correspondence between P and $\mathcal{R}$ is observed for real world networks such as the Austrian power grid, displayed in fig. 2(a), where the flow has been computed using the network flow model outlined above. The peaks in $\mathcal{R}$ and $\mathcal{P}$ clearly coincide as shown in fig. 2(b). Further examples of the correspondence are shown in fig. 2(c) and fig. 2(d), which show $\mathcal{R}$ and $\mathcal{P}$ for the IEEE 14 bus and 118 bus test networks [30,31]. Here the peak values of $\mathcal{P}$ are ~ 1.035, corresponding to a value of $\mathcal{R}$ indicating an average 0.4% increase in efficiency available to the whole system from reducing the capacity of a single edge; as this is a per edge value, it reveals a substantial amount of inefficiency across the network as a whole.

Key to what follows is that the maximum values of $\mathcal{P}$ and $\mathcal{R}$ occur at the same flow volume F. Determination of $\mathcal{R}$ is computationally onerous, requiring the evaluation of a convex optimisation problem for each of a network's edges, rendering it impractical for all but the smallest of networks. Evaluating $\mathcal{P}$ therefore provides a simple computational proxy for identifying regimes of relative redundancy, enabling very large networks of complex topology and composition to be investigated. The algorithm for computing $\mathcal{R}$ in a complex network is presented in the below.

Computation of $\mathcal{R}$. – Recall that $\mathcal{R}$ is defined as the mean relative increase in flow efficiency attainable by capping the capacity of an edge in the network. This requires computing the optimal amount by which each edge should be capped, which can be evaluated analytically for the network in fig. 1(a). However, $\mathcal{R}$ is not analytically tractable in the general case of complex networks with overlapping paths from sources to sinks; therefore, the method outlined in algorithm 1 is used.

This algorithm takes a graph $G = (\mathcal{V}, \mathcal{E}, c)$ comprising a set of nodes and edges, $\mathcal{V}$ and $\mathcal{E}$ respectively, together with a set of edge functions c, and compares the Nash flow volume on each edge to the GO flow volume on that edge in order to determine by how much its capacity should be

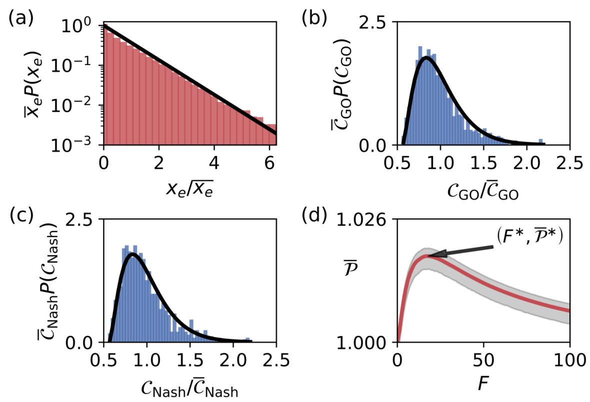

Fig. 3: (a) The Nash equilibrium edge power $x_e = c_e(f_e)f_e$ distribution in small world networks with $q = 0.1$, $k = 4$, $n = 32$, $n_s = n_d = 16$, $n_p = 0$. Distributions of (b) $\mathcal{C}_{\mathrm{GO}}$ and (c) $\mathcal{C}_{\mathrm{Nash}}$ in an ensemble of 1000 such networks with total flow volume $F = 20$. The solid lines are fitted shifted gamma distributions with shape parameter $\nu = 3.44$, scale parameter $\mu = 3.50$ and shift parameter $\sigma = 12.3$ in (b) and $\nu = 3.51$, $\mu = 3.51$ and $\sigma = 12.4$ in (c). (d) The mean Price of Anarchy $\overline{\mathcal{P}}$ as a function of F, with maximum at $(F^*, \overline{\mathcal{P}}^*)$. The shaded region indicates the 95% confidence interval, computed using the statistical bootstrapping method [28].

Algorithm 1 Compute $\mathcal{R}$

Input: A network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, c)$
Output: The redundancy measure $\mathcal{R}$ on $\mathcal{G}$

```
Compute the Nash and GO flows f_Nash and f_GO
C_Nash = Σ_{e∈E} c_e(f^e_Nash) f^e_Nash
for e ∈ E do
    if f^e_Nash > f^e_GO then
        Set an upper limit κ = f^e_GO on edge e
        Compute modified Nash flow f'_Nash
        C'_Nash = Σ_{e∈E} c_e(f'^e_Nash) f'^e_Nash
        R^e = (C_Nash − C'_Nash)/C_Nash
    else
        R^e = 0
    end if
end for
R = mean(R^e)
```

capped. A new Nash flow $\mathcal{C}'_{\mathrm{Nash}}$ is then computed after capping edge e, from which $\mathcal{R}^e$ is then computed. This process is repeated for all edges to obtain the mean $\mathcal{R} := \overline{\mathcal{R}^e}$. For some edges there may be no possible improvement in cost by removing capacity, in which case we set $\mathcal{R}^e = 0$. For examples of results using this method, see fig. 1(d) and fig. 2(b)-(d).

Dependence of $\mathcal{P}$ on flow and network composition. – We first consider networks whose source and sink nodes have homogeneous flow outputs and inputs respectively, given by the case where $\xi_v = 0$ for all v in eq.(2). For a total flow volume F, the dependencies of $\mathcal{P}$ on network structure and composition are obtained from an ensemble of 1000 such random small-world network realisations [32, 33]. These networks are parameterised by the rewiring probability q, initial degree k, and the number of nodes n, comprising n_s, n_d and n_p source, sink and passive nodes respectively, whose locations are randomly allocated. The edge cost coefficients α_e and β_e are both uniformly distributed random variables in the range [0,1]. At the microscopic scale in the network, fig. 3(a) shows that the individual edge costs are exponentially-distributed. Unsurprisingly, at the macroscopic scale the total Nash and GO costs (representing total power loss) are gamma-distributed with a probability density function $P(\mathcal{C}) = (\mathcal{C} - \sigma/\mu)^{\nu-1} \exp(-(\mathcal{C} - \sigma)/\mu)/\mu\Gamma(\nu)$, since they are formed from an ensemble of exponentially-distributed edge costs. This is shown in fig. 3(b),(c) and confirmed by Kolmogorov–Smirnov tests (see *supplementary material* for more detail). For each F, the mean of the resulting distribution of $\mathcal{P}$, denoted $\overline{\mathcal{P}}$, is shown in fig. 3(d). With increasing flow, $\overline{\mathcal{P}}$ rapidly rises to a maximum $\overline{\mathcal{P}}^*$ at F^*, before declining to unity. How the values of $\overline{\mathcal{P}}^*$ and F^* depend on the network configuration, defined by n_s, n_d and n_p is now considered.

The condition $n_s + n_d + n_p = n$ constrains the space of possible network configurations to a triangular-shaped simplex whose vertices touch one of the n_s, n_d, n_p axes, as depicted in fig. 4(a). The variation of $\overline{\mathcal{P}}^*$ and F^* for constant n are then projected onto this simplex, as shown in fig. 4(b),(c), respectively. The contours are symmetric about a line bisecting the simplex, corresponding to networks with $n_s = n_d$ and shown by section (i) in fig. 4(b). Along this line the value of $\overline{\mathcal{P}}^*$ decreases monotonically with increasing n_s, as shown by the plot in fig. 4(d). Section (ii) is a slice across the simplex at whose mid point $n_s = n_d$. $\overline{\mathcal{P}}^*$ increases monotonically as this point is approached from either direction, as shown in fig. 4(e), revealing that inefficiency and average edge redundancy are maximised as the number of source and sink nodes becomes equal. fig. 4(f) shows $\overline{\mathcal{P}}^* \sim a + bn_s^{-1/2}$ on section (iii), along which n_s increases (and n_p decreases) with $n_d = 1$. The morphology of the contours shown in fig. 4(b),(c) remains invariant with q, meaning that these results pertain to both small-world and random Poisson ($q > 0.6$) networks, as demonstrated in fig. 5. This invariant property also persists (*supplementary material*) when considering scale-free networks [34], whose topology is quite distinct from the small-world and Poisson classes.

In practice sources and sinks may be expected to have heterogeneous levels of output and input, such as an electrical grid containing a range of generators with different output capacities. To account for this, ξ_v in eq.(2) is now set to be a normally distributed random variable with mean 0 and variance 0.2. This represents a substantial amount of heterogeneity whilst typically still preserving the types of the nodes, and therefore the location on the simplex. Fig. 6 demonstrates this heterogeneity in en-

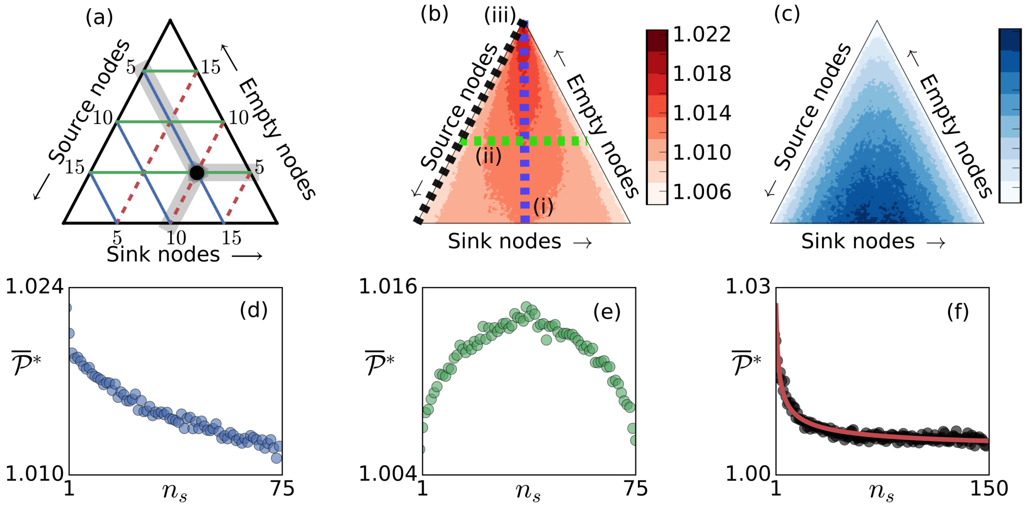

Fig. 4: (a) A sketch of the node configuration space simplex. The black dot represents a configuration of $n_s = 5$, $n_d = 10$ and $n_p = 5$. (b,c) $\overline{\mathcal{P}}^*$ and F^*, respectively, for ensemble of 500 small world networks, each with $n = 150$, $k = 4$ and $q = 0.1$ projected onto the simplex in (a). (d–f) $\overline{\mathcal{P}}^*$ as a function of n_s along the sections (i–iii) indicated in (b). In (f), the red line indicates the function $a + bn_s^{-1/2}$ with $a = 1.003$, $b = 0.024$.

sembles of small world networks and reveals that the key features of the simplex remain. In particular the highest values of $\overline{\mathcal{P}}^*$ are found on the centre line of the simplex where the numbers of sources and sinks are equal.

Whilst the morphology of the contours remains approximately invariant with network size (*supplementary material*), the values of $\overline{\mathcal{P}}^*$ and F^* do scale with network size. Fig. 7(a) shows that the maximum value of $\overline{\mathcal{P}}^*$ for small-world, Poisson and scale free networks saturates to a constant value for $n > 50$, whereas fig. 7(b) shows that F^* increases linearly with network size. These scaling results can be used in conjunction with fig. 4 to interrogate networks of arbitrary size.

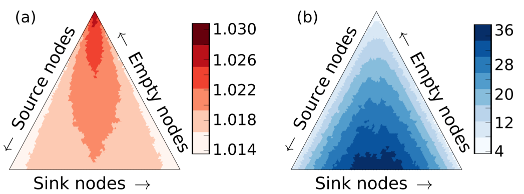

Fig. 5: (a) $\overline{\mathcal{P}}^*$ and (b) F^* for an ensemble of 500 random Poisson networks each with $n = 64$, generated by the Watts-Strogatz method [32] with $q = 0.6$ and $k = 4$.

The linear scaling shown in fig. 7(b) can be explained. F^* corresponds to a threshold beyond which the network flows adjust such that the two equilibrium costs begin to converge. To exceed the threshold the total flow must increase linearly because the expected density of flow decreases with increasing n.

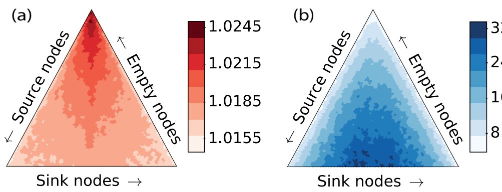

Fig. 6: (a) $\overline{\mathcal{P}}^*$ and (b) F^* for an ensemble of 500 random $n = 64$ small-world networks with q=0.1. In these networks ξ_v is a normally distributed random variable, with mean 0 and variance 0.2, inducing sources and sinks to have heterogeneous flow inputs and outputs.

Conclusion. – This Letter has investigated how inefficiency of flows occurring on different classes of random network, as gauged by the Price of Anarchy $\mathcal{P}$, is affected by the network structure and the function of its nodes. It has also established a correspondence between $\mathcal{P}$ and measures of network redundancy, an important consideration in addressing issues of network resilience and cost-effectiveness. This is primarily motivated by understanding properties associated with flows of current in electrical micro-grids, wherein nodes are either sources or sinks of current, or are passive conduits. Poisson, scale-free and small-world networks are used to establish the generality of the results with respect to network topology; this reveals a predictable dependence of $\mathcal{P}$ upon node composition for networks of arbitrary structure.

The simplex plots fig. 4(b) and (c) and their symmetry and invariance properties, when taken in conjunction

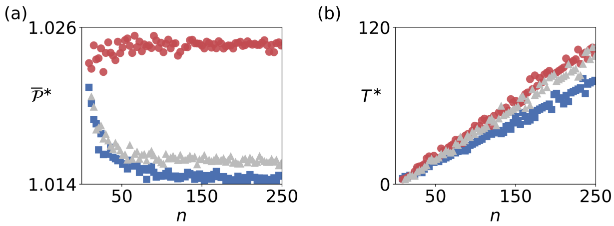

Fig. 7: (a) $\overline{\mathcal{P}}^*$ and (b) F^* as functions of n for small world networks (blue squares) with $q = 0.1$ and $k = 4$; Poisson networks (gray triangles) generated using the Watts-Strogatz method with $q = 0.6$ and $k = 4$ and scale-free networks (red circles), generated using the Barabási-Albert method [34]. All networks are chosen to have a node configuration $n_s = n_d = n_e$.

with the system size scalings shown in fig. 7, provide an operating space that defines maximal inefficiency and redundancy for an ensemble of networks with general topology and with variable node composition. With application to micro-grids, a given network's composition will change both diurnally and seasonally, traversing a trajectory through this configuration space. This path will depend on the nature of the sources of power and the load consumed by the sinks – features that will vary with population behaviors and the variable outputs from renewable power sources, for example. This information can be exploited to aid in the dynamic design and management of smart networks so as to constrain trajectories to preferred regions on the simplex. Insofar as redundancy is related to resilience [35–37], this aspect of the system's performance can be manipulated dynamically via the network's node type configuration and edge costing. A striking feature is that greatest values of inefficiency (or redundancy) occur when the number of sources and sinks are equal, as apparent in fig. 4(b) and fig. 5(a), a situation that is prevalent for small renewable energy networks where the numbers of generators and consumers are comparable. By contrast, the results show that a centralised electrical distribution grid comprising a few sources but many sinks has a low $\overline{\mathcal{P}}^*$, indicating it is both efficient and lacks redundancy. Equivalent plots can be constructed that are particular for an individual network's structure and composition with which its performance can be gauged.

These findings have established that even for simple linear edge functions, network topology and flow conservation laws are sufficient to induce inefficiency that depends predictably on the configuration of nodes. An interesting extension to this work would be the consideration of nonlinear cost functions, for which the values of $\mathcal{P}$ may be substantially larger [11, 14].

The inefficiency caused by redundancy is only one metric with which to assess performance and it is inefficient networks that will generally also be the most resilient to faults or attack. Redundancy may also give networks flexibility to operate in a variety of conditions; however, since inefficiency and redundancy coincide, as we show in our results, optimising a network's structure and composition purely for efficiency may result in a loss of useful redundancy. Hence in using the simplex to aid network design it is likely that options will be constrained to an operating space offering an acceptable efficiency-resilience trade-off.

* * *

OS acknowledges the support of the Leverhulme Trust via the "Modeling and Analytics for a Sustainable Society" doctoral training scheme.

REFERENCES

[1] Roughgarden T., *Selfish Routing and the Price of Anarchy* (The MIT Press) 2005.
[2] Nash J. F., *Proc Natl Acad Sci USA*, **36** (1950) 48.
[3] Nash J. F., *Ann Math*, **54** (1951) 286.
[4] Koutsoupias E. and Papadimitriou C., *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science* STACS'99 (Springer-Verlag, Berlin, Heidelberg) 1999 pp. 404–413.
[5] Cohen J. E., *Proc Natl Acad Sci USA*, **95** (1998) 9724.
[6] Fabrikant A., Luthra A., Maneva E., Papadimitriou C. H. and Shenker S., *Proceedings of the Twenty-second Annual Symposium on Principles of Distributed Computing* PODC '03 (ACM, New York, NY, USA) 2003 pp. 347–351.
[7] Andelman N., Feldman M. and Mansour Y., *Games Econ Behav*, **65** (2009) 289 .
[8] Knight V. A. and Harper P. R., *Eur J Oper Res*, **230** (2013) 122 .
[9] Perakis G. and Roels G., *Management Science*, **53** (2007) 1249.
[10] Correa J. R., Schulz A. S. and Stier-Moses N. E., *Games and Economic Behavior*, **64** (2008) 457 .
[11] Youn H., Gastner M. T. and Jeong H., *Phys Rev Lett*, **101** (2008) 128701.
[12] Cohen J. E. and Horowitz P., *Nature*, **352** (1991) 699-701.
[13] Roughgarden T. and Tardos E., *J ACM*, **49** (2002) 236.
[14] Roughgarden T., *J Comput Syst Sci*, **67** (2003) 341 .
[15] Skinner B., *Phys Rev E*, **91** (2015) 052126.
[16] Rose A., O'Dea R. and Hopcraft K. I., *Phys Rev E*, **94** (2016) 032315.
[17] Braess D., Nagurney A. and Wakolbinger T., *Transportation Science*, **39** (2005) 446.
[18] Roughgarden T., *J. Comput. Syst. Sci.*, **72** (2006) 922.
[19] Steinberg R. and Zangwill W. I., *Transportation Science*, **17** (1983) 301.
[20] Witthaut D. and Timme M., *The European Physical Journal B*, **86** (2013) 377.
[21] Witthaut D. and Timme M., *New Journal of Physics*, **14** (2012) 083036.
[22] Coletta T. and Jacquod P., *Phys. Rev. E*, **93** (2016) 032222.
[23] Motter A. E. and Timme M., *Annual Review of Condensed Matter Physics*, **9** (2018) 463.

[24] Brummitt C. D., Hines P. D. H., Dobson I., Moore C. and D'Souza R. M., *Proc Natl Acad Sci USA*, **110** (2013) 12159.
[25] Albert R. and Barabási A.-L., *Rev Mod Phys*, **74** (2002) 47.
[26] Bertsekas D., *Nonlinear Programming* (Athena Scientific) 1995.
[27] Pigou A. C., *The economics of welfare* (Macmillan London) 1920.
[28] Efron B. and Tibshirani R., *Statist. Sci.*, **1** (1986) 54.
[29] Retrieved on 17/05/2019 from: https://www.apg.at/en/Stromnetz
[30] Retrieved on 17/05/2019 from: https://icseg.iti.illinois.edu/ieee-14-bus-system/
[31] Retrieved on 17/05/2019 from: https://icseg.iti.illinois.edu/ieee-118-bus-system/
[32] Watts D. J. and Strogatz S. H., *Nature*, **393** (1998) 440 EP .
[33] Amaral L. A. N., Scala A., Barthélémy M. and Stanley H. E., *Proc Natl Acad Sci USA*, **97** (2000) 11149.
[34] Barabasi A.-L. and Albert R., *Science*, **286** (1999) 509.
[35] Halu A., Scala A., Khiyami A. and González M. C., *Sci Adv*, **2** (2016) .
[36] Quattrociocchi W., Caldarelli G. and Scala A., *PLOS ONE*, **9** (2014) 1.
[37] Corson F., *Phys Rev Lett*, **104** (2010) 048703.