

Shape and Performance of Fastest Paths over Networks with Interacting Selfish Agents

Marco Cogoni, Giovanni Busonera, and Enrico Gobbetti
*CRS4 - Center for Advanced Studies, Research and Development in Sardinia and
Ex Distilleria - Via Ampere 2, 09134 Cagliari (CA) Italy**

We study the evolution of the fastest paths in transportation networks under increasing congestion, modeled as a linear decrease in edge travel speed with density due to interactions among selfish agents. Moving from the common edge-based to a path-based analysis, we examine the fastest directed routes connecting random origin-destination pairs as traffic grows, characterizing their shape through effective length, maximum detour, and area under the curve, and their performance through a novel metric measuring how fast and how far an agent travels toward its destination. The entire network is characterized by analyzing the performance metric's distribution across uniformly distributed paths. The study covers both random planar networks with controlled characteristics and real urban networks of major cities. The low-density network regime, in which an initial smooth performance degradation is observed up to a critical traffic volume, is followed by the emergence of complex patterns of spatially heterogeneous slowdowns as traffic increases, rapidly leading to disjoint subnetworks. The failure of a few edges leads to a catastrophic decrease in the network performance. The fastest paths for all cities show a peak for detour and inness (and their variance) in the proximity of the critical traffic level, defined as the flex of the rejected path ratio curve. Inness generally shows a slight attraction by city centers on paths for light traffic, but this reverses to strong repulsion during congestion. We exploit path performance to uncover an asymmetric behavior of different regions of the networks when acting as origins or destinations. Finally, the Gini coefficient is used to study the unequal effects of path performance degradation with traffic.

I. INTRODUCTION

Network transport performance and robustness are usually estimated using static measures depending on topological and geometrical factors that characterize the network in its empty state, but cannot grasp the complex structural modifications that happen when the network is subject to growing levels of traffic [1–3]. Network robustness is often studied by removing random elements such as nodes and edges or varying their attributes to highlight weak points and bottlenecks within the system [4]. While very useful for evaluating abstract network geometries, this approach does not consider any specific process happening over the network. It turns out, however, that details of the transport process may greatly impact how the system reacts to perturbations. In particular, perturbations are typically non-random and spatially correlated, and knowing the specific process allows the development of better testing procedures [5, 6]. In this context, percolation theory is sometimes used to model transport disruption [2, 7], and it can be used to estimate the fraction of malfunctioning edges (or nodes) leading the network to break up into several disconnected clusters [8]. This approach has been recently applied to traffic in urban networks by using the vehicle speed on each road (edge) as a predictor of transport collapse to find the threshold speed above which fragmentation appears for different traffic regimes [9]. Percolation, however, ignores the details of the transport process altogether, modeling only its effects.

Estimating whole-network performance using the distribution of edge efficiencies may give a biased picture depending on the specific transport process in the system. E.g., when the transports happen between *Origins* and *Destinations* (OD) connected via the shortest (or fastest) paths, the chosen edges are spatially correlated, so estimating the global transport efficiency by averaging over the edges is not accurate. Different requirements on path routing and OD spatial distribution may lead to entirely different network usage patterns, thus translating to different expected performances. Transport processes are found everywhere, for instance in urban networks, whose traffic patterns have been modified by the emergence of personal navigation tools providing the fastest routes for their users in real time [10].

A non-local measure widely employed to estimate network performance at the edge level is the *Betweenness Centrality* (BC). BC and other centrality measures have been used to predict which edges (and nodes) are typically subject to the highest traffic, but the correlation with real traffic suffers in the medium and high-density regimes [11, 12]. Edge usage obtained from BC, in fact, typically mimics a low-density state of the network that usually happens with very small (compared to network geometry) or very fast (with respect to congestion buildup timescales) agents [13].

Recently, standard BC has been improved by introducing a model to simulate urban networks with different traffic levels and considering the strong interaction effects observed in different conditions to obtain a better measure of edge congestion [9]. While the approach's theoretical foundations are transferable to other contexts [14–17], urban networks have specific connectivity character-

* {marco.cogoni, giovanni.busonera, enrico.gobbetti}@crs4.it

istics that have led to the emergence of specific theoretical models and studies. In particular, they belong to the special class of (almost) planar graphs [1, 18] whose topology is constrained by the geographical embedding. This severely hinders their long-range connectivity and limits their maximum node degree [19]. The specificity of urban networks has been widely studied [10, 17, 20, 21], both in terms of their growth over time and of their complex dynamics under different traffic levels.

This work proposes shifting the measuring process from an edge-based to a path-based network performance metric. The basic idea relies on the fact that having a large fraction of network components that are still efficient does not guarantee sufficient global performance if those components are seldom used for the assigned task. This is especially crucial for transport processes such as urban networks, where the congestion patterns are shaped by interacting traveling agents that react with different routing strategies as traffic grows.

We first study the scaling behavior of fastest-path length and detour over two orders of magnitude of geodesic distance for a set of large, densely-populated, metropolitan areas and for two kinds of random planar graphs. Critical exponents are then computed for growing network congestion and plotted on a plane to better understand how they evolve. While critical exponents give an idea of how the fastest paths change shape, to better characterize their morphology over a broad range of traffic, we also analyze inness [22] to check how the supposedly attractive force exerted by city centers behaves. Moving beyond path morphology we discuss path performance evolution and its distribution. Path performance turns out to be useful not only in better evaluating whole-network efficiency, but in investigating the asymmetry of different parts of the graph when acting as sources or destinations of traffic. Finally, we compute the Gini coefficient of the performance distribution to quantify inequality over paths for all cities and show that it can be used as a sensitive measure to probe the critical traffic level of a network.

II. METHODS

A. Interaction model

In real urban networks, the average travel time is of the same order of magnitude as the typical timescale of congestion buildup, which has been measured to be about one hour in major cities [20]. The strength of the interaction among vehicles depends on the local traffic density given by the probability of vehicle coexistence on the same edge during rush hours. This allows for estimating the cumulative traffic seen on the roads during that finite time window. The interaction model used in this work does not incorporate full vehicle dynamics but takes into account the contribution to the traffic of each vehicle added to the network at a road segment level by updat-

ing the expected travel times at each step. An earlier work [9] has shown how this model, though simplified, can describe qualitative and quantitative properties of the network loading process of synthetic and real urban networks. We summarize here the aspects of the model relevant to this work and refer the reader to the original publication for further model details [9].

Time is not modeled explicitly, but the interval between the empty network and congestion, which represents the rush-hour period, is divided into a sequence of vehicle creation steps, which may be added one at a time or in batches. At each step, the fastest route between O and D is computed with complete knowledge of the network state, and the vehicle density (i.e., speed) on each affected edge is updated. In this model, vehicles are never removed from the network, even the ones traveling for very short paths, though their impact on the network will be proportional to the ratio of expected travel time on the affected edges and the chosen simulation duration (τ). The relation between speed and density on a single edge follows the single regime Greenshields model [1, 23], for which speed starts as free flow on an empty road, decreasing linearly to zero with maximum density.

We simulate the network evolution, as observed by travelers, while the traffic V increases from zero up to almost complete gridlock, signaled by the vanishing probability of adding new paths not containing congested edges. The traffic network is modeled as a directed, weighted graph whose edges represent road segments between adjacent intersections (nodes) and possess three constant features: physical length l , maximum speed, and number of lanes. Nodes, on the other hand, are featureless. Travel time t on each edge is defined by the ratio of its current speed and length.

The network traffic grows incrementally by activating one new path at each simulation step, to reach the desired target volume at time τ (the end of the simulation). Thus, simulation steps may be interpreted both as the current number of added paths and as a temporal marker within the sequence of OD pairs randomly generated for each simulation. O and D are chosen uniformly over the network nodes.

Intuitively, the occupancy factor induced by a vehicle over an edge is proportional to the time the vehicle is supposed to spend on it (as a fraction of τ), as forecast at its departure, and the sum over the whole path will be equal to unity (certainty of finding the vehicle on the path during τ) only when the travel time is greater than τ . Since the initial vehicles find a nearly empty network, their fastest paths and travel times are virtually equivalent to the non-interacting case. With rising traffic, however, edges fill up and the previous fastest paths will disappear and less-used roads and residential neighborhoods will be chosen. Some edges will eventually reach maximum density and become congested. If the fastest route from O to D comprises a congested edge (i.e., the network is disconnected) we still choose to add the initial part of the path, but skip all remaining edges from the

first congested one.

B. Path shape analysis

We characterize the morphology of a path from O to D (with $s = |OD|$ being the geodesic distance) by looking at two geometrical properties:

- Detour $d = \max_{i \in p} h_{\perp}^i$: maximum distance between the path and the geodesic line;
- Inness $I = \sum_{i \in p} h_{\perp}^i h_{\parallel}^i$: signed area between the path and the geodesic line;

where h_{\parallel}^i stands for the scalar product of the i -th directed edge with \overrightarrow{OD} , while h_{\perp}^i is the edge distance from \overrightarrow{OD} . Fig. 1 offers a visual definition of these objects.




FIG. 1. Between O and D we construct: a geodesic s (dashed red), the fastest path p^e on the empty graph (blue), and the fastest path p for the loaded graph. Since p contains a dysfunctional edge, p^* is the longest allowed path.

The sign of the inness is taken to be positive for paths traveling toward the center of the map, negative otherwise: Thus, it will be positive whenever a “centripetal force” [22] exists. Normalized I is obtained by division with the square built on the geodesic line: $\tilde{I} = I/s^2$ and normalized detour as $\tilde{d} = d/s$.

C. Slowdown, Completeness and Performance Index

Given the network in some generic state and an OD pair, the fastest path p (composed of a sequence of edges of length l_i) is computed. A special case of p is the fastest path as computed over the network in its empty state: p^e . We define its Completeness factor (C) as the ratio between the distance $|p^*|$ that can actually be traveled along p , and the whole path length $|p|$:

$$C^{OD} = \frac{|p^*|}{|p|} = \frac{\sum_{i \in p^*} l_i}{\sum_{i \in p} l_i}.$$

The admissible traveling distance over the path is either limited by the presence of unavoidable dysfunctional edges (i.e., the network is fragmented) or because τ would be exceeded. In both cases, we load the network only for the reachable part of the path p^* .

The Slowdown factor (S) is ideally defined as the ratio between the travel time from O to D on the fastest path over the empty network $T_{p^e} = \sum_{i \in p^e} t_i^e$, and the one necessary on the fastest path available (typically a different one) on a congested system $T_p = \sum_{i \in p} t_i$. This simple definition must be slightly modified to take into account those paths that cannot be completed either because of disconnected O and D (hard condition) or because travel time would be longer than the chosen τ (soft condition). In these situations, to be meaningful, the slowdown computation is limited to the portion of the path that is possible to travel over the congested network by rescaling the time on the empty network by C^{OD} :

$$S^{OD} = \frac{T_{p^e}^{OD} \cdot C^{OD}}{T_p^{OD}}.$$

T_{p^e} is the total travel time along the shortest path p^e between O and D when the graph is empty. T_{p^*} is the travel time of the OD path until the dysfunctional edge in the case of the loaded graph. $T_p > \tau$ values are clipped.

Finally, we introduce the *Performance Index* (P) for a path connecting O and D as the product of S^{OD} and C^{OD} for any network state:

$$P^{OD} = S^{OD} \cdot C^{OD}. \quad (1)$$

In a low congestion regime, P will simply follow the slowdown behavior, whereas for high traffic volumes, the completion ratio will often dominate, with several paths containing dysfunctional edges preventing the transport from reaching its destination. The P index is a non-local performance metric that evaluates the efficiency of the transport over the paths as opposed to standard metrics that focus on the performance of the single network constituents such as edges and nodes. Transport performance is heavily influenced by the kind of process taking place over the network and induces a specific spatial correlation on edge usage that no uniform averaging over the edges can approximate. Our choice of the C factor is the only part of this definition that is tailored for vehicular traffic since it assumes that a partially completed transport is still valuable.

D. Non-uniform path performance degradation

The Gini coefficient has been used in recent years to describe structural properties of networks, such as the inequality of the degree and BC distributions [24], and to extract structural profiles of urban networks [18]. We compute the Gini coefficient of the P distribution to evaluate inequalities in path performance degradation. To measure inequality in a set of quantities P_i we define:

$$G = \frac{\sum_i \sum_j^n |P_i - P_j|'}{2n \sum_i^n P_i}$$

The Gini coefficient tends to zero (perfect equality) when most paths have similar P values; on the other side of

the spectrum, maximum inequality among P values is signaled by $G \rightarrow 1$: i.e., only a few paths maintain acceptable performance while the rest is dysfunctional.

III. RESULTS AND DISCUSSIONS

We selected eight large metropolitan areas: Beijing, Berlin, Las Vegas, London, Los Angeles, Madrid, Paris, and Rome to represent very different topologies, mainly stemming from different geographical locations and historical evolution processes. The vehicular transport network of the areas around city centers (20 km radius except for Rome and Madrid, with a radius of 12 and 15 km, respectively) were downloaded from OpenStreetMap (OSM) retaining the information about edge length, maximum speed, and number of lanes. The number of edges of the corresponding graphs is in the range 1.0×10^5 (Rome) to 5.5×10^5 (London). The maximum traffic load we generated for all cities was limited to 2.0×10^6 OD pairs, enough to drive all cities into a deeply congested state. The maximum allowed time for a path to be completed was set to $\tau = 3600$ s. All vehicles travel at the maximum speed allowed on the edges, depending on the state of the simulation at a specific time. In contrast to the simulations on the same graphs performed for a previous publication [9], we do not perturb the values with Gaussian noise to simplify the analysis¹.

A. Path shape analysis

We perform an analysis of the shape of the paths modeled after previous works on spatial networks [22, 25]. Each path is characterized by the following geometrical properties and their scaling with the geodesic OD distance:

- path travel time T and length $|p|$;
- *detour* (d): maximum distance along the path from the geodesic line between O and D;
- *inness* (I): the area comprised between the path and the geodesic, considered as positive when leaning toward the center of the map.

We study how the above properties change as the network is subject to increasing traffic levels. The scaling behavior with the OD distance is studied in a log-log space by performing a linear fit to extract the “critical” exponents. A strict definition of these exponents holds only when considering infinite networks, but the idea is useful to characterize real finite networks [20]. We focus on scaling exponents χ and ξ , for path length standard

deviation from the mean $\sigma(|p|) \sim s^\chi$, and for the wandering $\mathbb{E}(|p|) \sim |s|^\xi$, respectively. We also compute the exponent m associated with path travel time T scaling (that scales as s^m). We denote the average of a quantity by $\mathbb{E}(\cdot)$ and $a \sim b$ means a converges to Cb with C a constant independent of OD, as $s \rightarrow \infty$.

1. Critical exponents in the low-congestion regime

Path travel time scaling (m) is well approximated by a linear fit in the log-log space for all cities except for Beijing where the curve does not follow a power law. The linear fit holds up only up to ~ 3 km for low traffic, while linearity is regained deep into the congested regime. The results of the linear fits performed with increasing congestion (below the critical level) for real cities and for two random planar graphs (Gabriel and k-Nearest Neighbors (K-NN)) are reported in tab. I.

city	χ	ξ	m
Beijing	0.10	0.55	0.73
Berlin	0.37	0.71	0.76
Las Vegas	0.30	0.72	0.75
London	0.36	0.69	0.76
Los Angeles	0.30	0.78	0.78
Madrid	0.23	0.74	0.72
Paris	0.33	0.74	0.74
Roma	0.22	0.65	0.76
Gabriel	0.40	0.70	1.00
k-NN	0.20	0.60	1.00

TABLE I. Critical exponents χ , ξ , and m computed for all cities with no traffic and for the undirected Gabriel (and k-NN) graph ($\rho = 10^{-4}$). Linear fits, and their complete evolution with growing traffic, are shown in figs. S17-S24




FIG. 2. Critical exponents χ and ξ for all cities (and for the Gabriel and k-NN graphs) for traffic levels up to 250k vehicles.

¹ We experimentally verified that results are not affected in relevant ways.

The values of χ and ξ shown in tab. I are drawn for all cities in fig. 2 as the smallest circles of each city and represent the empty network state that increases in size up to 250k vehicles. The Gabriel (defined by excluded region) and the k-NN (defined by node proximity) graphs are also included as a reference in their undirected versions with densities comparable to real cities (mean edge length about 80 m). London shows a significant variation of the fitted exponents already with the addition of a few thousand vehicles (because its transition to congestion happens earlier at ~ 0.3 M vehicles) while other cities are more stable in this low traffic regime with about half the drift: This drift involves in all cases a decrease of both ξ and χ for all cities but London and, to a lesser extent, Berlin (increase of χ and later ξ inversion), and Rome (almost constant). London and Berlin are very close to Gabriel, and Rome ends very near the k-NN. Beijing has the lowest values for both exponents ($\chi \sim 0.1$ and $\xi \sim 0.55$) while Los Angeles, Las Vegas, Paris, and Madrid have intermediate $\chi \sim 0.3$ and the highest $\xi \sim 0.75$ initial values. A coherent comparison of real cities with synthetic graphs involves directional edges for the Gabriel and k-NN and some form of sparsification (to remove perfect bidirectionality) and will be the subject of future work. Critical exponents within the congested phase are not discussed since the fits are noisy (especially χ) with large error bars beyond the transition: their full behavior is visible in figs. S17-S24, though.

2. Morphological changes and central force effects

The ξ and χ exponents describe the scaling of the detour and path-length variability for a fixed traffic volume, but we are interested in studying how their distributions change at all traffic levels. For the city of London, $\mathcal{P}(d)$ and $\mathcal{P}(\sigma(|p|))$ are shown in fig. 3 where the color indicates the number of paths in each bin. The inness and detour distributions are wide and non-Gaussian in general, but we also report their average values in fig. 4 to better visualize where the majority of paths is located.

For low traffic levels (below 0.2×10^6 vehicles) the distribution is limited to a maximum of about 10 km detours over the whole network, with most values below 5 km. Increasing traffic toward the transition (red vertical dashed line at about 0.3×10^6 vehicles) the distribution suddenly expands to $d > 40$ km for a relevant number of paths that are trying to avoid preferential (faster, shorter, multi-lane) roads that have become congested. This transition traffic volume (clearly specific for each network) coincides also with the position of the flex in the curve of the hard-rejected paths in fig. 5. The hard reject ratio seems to be a good predictor of this abrupt change in the network transportation performance, while the soft reject curve peaks before, and in the London case, reaches about 50% of the paths being soft-rejected for about 0.25×10^6 vehicles. While the soft rejection depends on the size of a city and the choice of $\tau = 3600$ s, the hard rejection

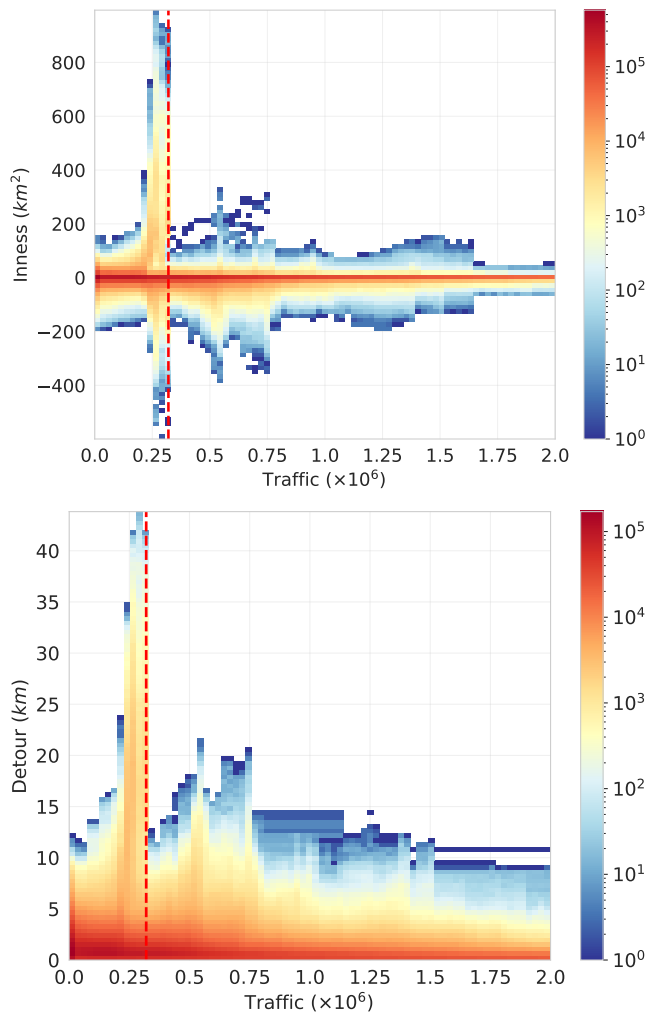


FIG. 3. The full distributions of inness and detour for London. The logarithmic color bar represents how many (#) paths fall in each bin, with blue for $\# < 10$ and red for $\# > 10^2$. The congestion transition is shown by vertical dashed red lines. Each column contains $\sim 5 \times 10^5$ vehicles.

signals the break-up of the network into multiple disconnected islands. In this case, it is interesting to note that having 50% of hard-rejects at the transition (see fig. 5) roughly tells that the bridges over the Thames have collapsed as will be confirmed in the following. The number of edges needed for this to happen is very small (typically $< 0.1\%$, see figs. S1-S4(bottom)) since the first edges to reach maximum saturation are the most desirable ones, on average. Thus, building resilient cities to better exploit their maximum traffic volume potential, naturally involves having multiple available paths among adjacent neighborhoods (that are usually mesh-like and well-connected) avoiding lone bridges that tend to saturate early regardless of their capacity. Several small bridges are usually more resistant than a single one with equivalent capacity because spatial multiplicity tends to distribute traffic over larger areas, better using the avail-




FIG. 4. Mean values of inness and detour for London. Small colored circles are reference traffic volumes used later. The congestion transition is shown by vertical dashed blue lines. Colored circles following the “gnuplot” colormap serve as a guide for traffic levels later used in fig. 7.

able potential road volume.

While detour d highlights the maximum distance of the path from the geodesic, another way to quantify a shape change in a path is to compute the whole area between it and the geodesic. Since the sign is positive-definite for paths leaning toward the map center, we can measure the effects of the central force due to topology or congestion of the city center [22]. In fig. 4(top) we see how the inness I evolves with traffic intensity: For an empty network, the mean value is very small and slightly negative (centrifugal). Increasing traffic, the central force disappears to later become centripetal, then exploding in a huge peak just before the transition for which I reaches a mean value of 40 km^2 , but some paths are jolted to the other side of the network (across the center, thus the inness positivity, but an effectively repulsive behavior nonetheless) reaching maximal values of more than $I \sim 800 \text{ km}^2$. The transition coincides with the maximum mobility (flatter cost-of-transport surface [26]) in the route choice with paths spanning the entire network in their selfish struggle to finish within τ (or finishing at all). The peaking of d and I happens exactly at the same traffic level, just before the transition. After this short-lived transitional regime (involving the addition of about 0.1M vehicles), we observe an abrupt decrease of I and a crossover to moderately negative values: in this regime, the network center becomes repulsive and stays this way up to very high congestion. Continuing to increase traffic, it is evident that the magnitude of the areas shrinks with $I \rightarrow 0$. This is due essentially to a large portion of proposed paths being hard-rejected as seen in fig. 5. The longest paths are rejected first, then only the shortest ones can be added to the network without O and D

belonging to different network islands. The same fate happens to detour with $d \rightarrow 0$ for the same reasons. Detour and inness distributions for all cities are visible in figs. S5-S8 with similar results, but generally showing a slightly attractive central force for all cities except for Rome in which the center strongly repels all paths but the shortest ones. Los Angeles and Las Vegas have, on the other hand, fairly attractive centers when empty.

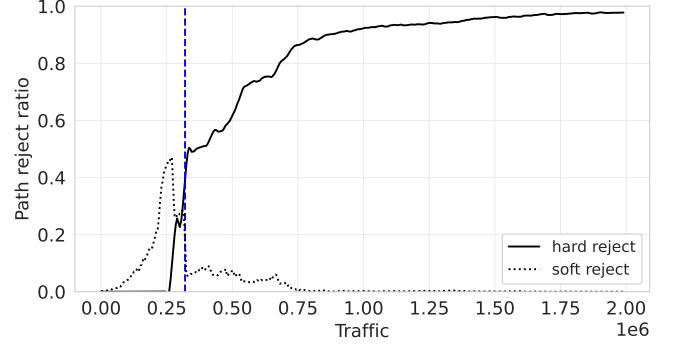


FIG. 5. Path reject ratio: solid black line are the hard-rejects (no path without dysfunctional edges available), dotted line are the soft-rejects whose travel time exceeds τ and can only complete part of their transport.

A normalized inness $\tilde{I} = I/s^2$ (divided by the squared geodesic distance) allows us to directly compare the effects of the central force on paths of different lengths at several traffic levels. We study \tilde{I} for several angular separations and distances of ODs from the center of the map to gather a better understanding of path shape dependence for specific OD configurations and traffic levels. In fig. 6 we plot \tilde{I} at four traffic levels for our reference city of London: Empty, 320k (the flex position), 475k, and 975k vehicles. With an empty network (left plot), the average \tilde{I} is very small for all path configurations, thus no central force effects are detected at any distance from the center in this case. Approaching the transition traffic level (second plot), the situation suddenly changes with the most dramatic effects seen at small angular (α) separations (5 km, blue line): for $\alpha \in (30^\circ, 90^\circ)$ the repulsive effect is strongest, with the average path being pushed away from the center to form an area of about half the OD square with the geodesic $\tilde{I} \sim -0.5$. This effect shows a chiral preference since for $\alpha \in (270^\circ, 330^\circ)$ (equivalent to the previous separations, but counterclockwise) the effect is much smaller with $\tilde{I} \sim -0.2$, but still repulsive. Error bars are very wide for all α except for 300° and 330° (30° and 60° CCW). For a larger radius (orange and green) the effect is weaker, but still repulsive for all angles and the error bars get smaller. At 20 km (red), $\tilde{I} \sim 0.0$ except for very small CW angles. For traffic levels beyond the transition (i.e., 475k and 975k vehicles) distributions become very narrow with error bars within symbol size and a small repulsive effect remains only for 5 km radius and small angle CCW paths. The

results regarding the empty networks, the only ones we can directly compare with our city selection, are quite different than the average ones published by Lee et al. [22]: Figs S9-S16 contain the same plots for the other cities.

To get a better idea of what happens to the fastest paths while increasing traffic, we show two examples in fig. 7 for the Greater London area that share similar OD separations but differ in orientation: In the left panel, the vehicle starts from south-west (red circle) to reach the east side (red diamond). Both O and D are south of the River Thames and the wide black path is the fastest path in the empty state: in this case, it is attracted toward the center with respect to the geodesic (dotted gray line). As traffic grows (violet to red) the force becomes strongly repulsive, peaking at about thrice the transition traffic when the red path is pushed toward the network’s southeast corner. Beyond this traffic level, no path exists for this OD and only a very small part of the journey (orange and yellow) can be traveled: this is expected from fig. 5 since about 85% of paths are rejected. On the right panel, we consider a North-South orientation with O and D on opposite sides of the river: the empty-network fastest path almost follows the geodesic (black), but gets repelled from the center at the transition traffic (violet). Beyond the transition, no available full paths exist (even the purple path stops very early). This breakup happens with three times less traffic than before because the river acts as the weakest part of the urban network connectivity, as expected.

B. Slowdown, Completeness and Performance Index

In previous sections, we saw how the shape of fastest paths depend on the congestion state of the underlying network and on the OD positions with respect to the map center. We now turn our attention to how shape impacts path performance at increasing traffic levels. Our simple traffic model linearly connects edge density to the slowdown experienced by vehicles on average. Thus, the most obvious effect that happens to vehicles when increasing the traffic level starting from an empty network is a degradation of travel times. This slowdown appears first on the most desirable roads (high speed, multi-lane, few traffic lights) since everybody chooses them as a first choice, especially for longer distances. As congestion starts to cripple the performance of these main roads, the fastest paths available necessarily switch to lesser-used (and capable) infrastructures. In fig. 8(top) we see how the population of paths connecting uniformly distributed OD pairs all over the network experiences slower speeds as the traffic volume increases. With zero traffic all paths are concentrated in the top left corner ($S \sim 1$) meaning maximum speed; as traffic grows, some paths slowdown and the distribution smears over wider values, with some OD at half their empty speed already with 100k vehicles, but most of them (average S of vertical bins shown

by horizontal black bars) maintain values > 0.8 . Right before the transition (red vertical line), the slowdown distribution widens spanning the whole $[0, 1]$ range with an average value of ~ 0.6 , with a non-negligible fraction of paths having one-tenth of their empty speed. A complex structure of oscillations for the S variance is evident from 0.2M to 0.8M happening in phase with further network breakups also visible in fig. 5: we can say that each “sub-transition” (with its smaller sigmoid/flex pattern) induces a new broadening of S and smaller peaks for detour and inness values (fig. 4). Slowdown then stabilizes at ~ 0.4 with a wide distribution until the maximum load (2M vehicles): nothing dramatic happens in this regime since the addition of longer paths is prevented, so almost only very short and partial paths are added. This fact brings us to the next quantity of interest to characterize travel efficiency: Completeness, which measures, for each OD pair, how far it is possible to travel along the fastest path. In urban networks, it makes sense to take into account this factor because between being able to travel 90% or 50% of the path results in a very different outcome from a traveler’s perspective. In other contexts, a non-linear behavior could be better suited, with the extreme case of all or nothing, such as in digital communication networks, where a data packet reaching 99% of its path is indistinguishable from one stopping near its origin. Even brewing espresso coffee would not gain much from water paths not fully reaching the lower side of the filter.

In this urban network context, for the city of London, we see that \bar{C} (black horizontal lines) drops from 1.0 steeply after a short plateau, and just before the transition, already one-third of the paths cannot reach its destination, either because $T_p > \tau$ or because only dysfunctional paths exist. The thin stripe of red bins at the top of fig. 8(middle), indicating completed paths, progressively becomes fainter and most paths prematurely stop after 1M.

In this context of urban transport makes sense to multiply S and C to obtain an overall estimate of how fast, and how far each path was able to go on its OD journey. We call this Performance Index P and we can see that its distribution is modulated by the S and C factors above it in fig. 8: The \bar{P} decrease is almost exponential with traffic and shows some non-trivial structures near the transition due to the succession of breakups of different parts of the city network. Beyond the transition, paths worsen their P from 0.4 to less than 0.1. A similar trend emerges for the other cities: see figs S33-S36. This behavior demonstrates that once we reach the vicinity of the transition, the overall performance of a realistic transport phenomenon deteriorates very quickly with small traffic variations: we can say that the network has a large susceptibility in this regime that we could define as the derivative of P with respect to traffic volume: $\frac{\partial P}{\partial V}$.

It is very useful to take the same P distribution for all paths at specific traffic levels (four vertical slices of fig. 8(bottom)) and map each path’s performance onto




FIG. 6. Normalized inness for four traffic levels (the second plot corresponds to the transition traffic level). Vertical dotted line divides CW from CCW paths: 330° means 30° CCW

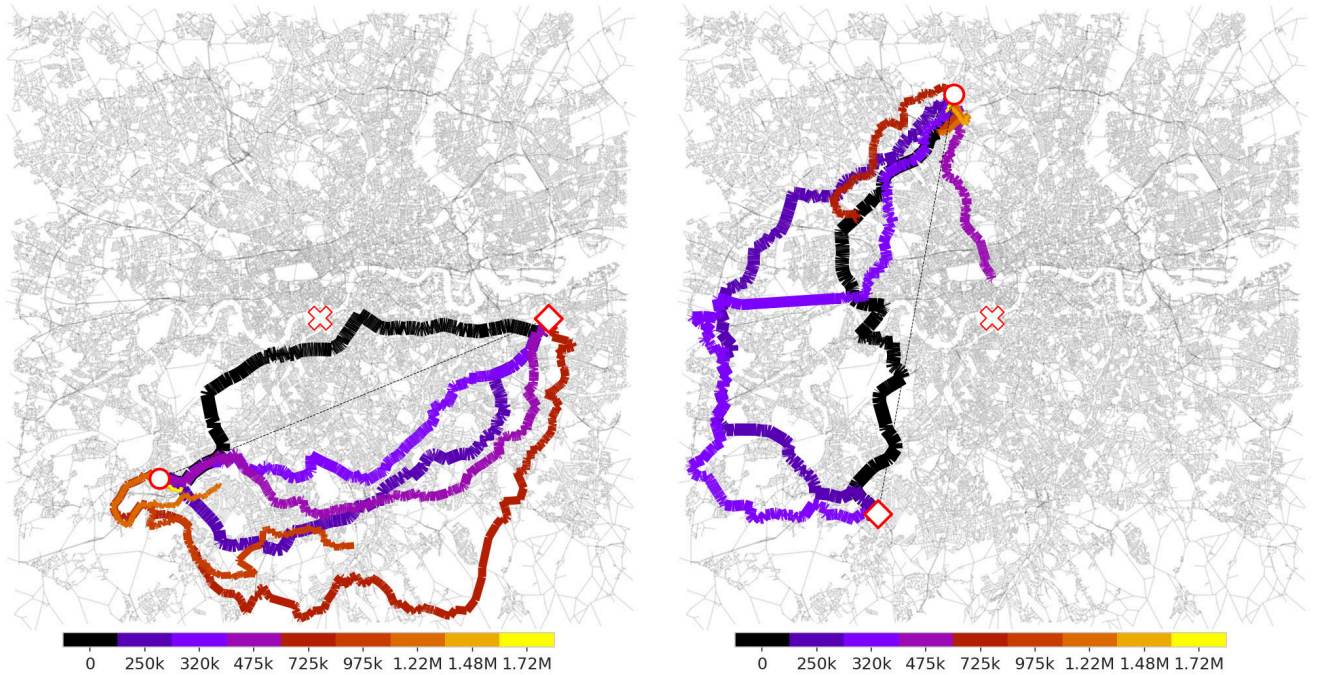


FIG. 7. In both panels: Fastest paths between same OD (both at 15 km radius) at different traffic load. (The circle and diamond symbols refer to Origin and Destination respectively, whereas the cross stands for the center of the map). On the left O and D are both below the river, while on the right the river must be crossed. The color bar shows the traffic level for each path. More angular configurations for all cities are visible in fig. S25-S32.

its origin fig. 9(top row) and to its destination fig. 9(bottom row) on the topographic chart. In this case, the color of a map location represents the average P of all paths leaving from (and going everywhere else) or arriving there from all other places. The left column shows the lowest traffic (50% of the transition flex) for both O (top) and D (bottom). The maps are almost unmarked in this case since most paths perform very well with $P > 0.7$. Moving toward higher traffic (95% of flex), P decreases fast and both O and D maps show different areas from which paths are starting (top) or arriving (bottom), that perform rather poorly (yellow ~ 0.5 and red ~ 0.3). Paths originating along the Thames (especially along the south

bank) seem to be the most impacted in this regime operating below 30% of their peak; other areas are still mostly functional. The situation about paths ending south of the Thames, on the other hand, already shows very poor $P < 0.3$. Increasing traffic just by 5% to reach the transition, we observe a slight improvement in P for the origins, especially the northern area, while the south bank of the Thames gets even worse. The D map at the flex seems to improve overall, but this is due to several paths not reaching their destination ($C < 1$), but with a good S along p^* , so P improves a little overall. Adding 25% more traffic to the city produces a catastrophic result on both maps with $P < 0.3$ almost everywhere. These




FIG. 8. Distributions (counts of paths per bin) of S , C and P for London for increasing traffic. The red line is the transition and the black curve is the average for each traffic bin (containing 25k vehicles). See figs. S33-S36 for the other cities.

P maps show mainly which parts of the network tend to break first, as already seen in previous works [8, 9], but also how certain areas perform differently when acting as traffic sources or sinks: Some places remain easily reachable during congestion, but, at the same time, leaving them to reach any other destination can be very hard. This asymmetry is partly related to what happens

in fig. 6 where CW and CCW paths at the same distance from the center could have very different d and I , especially at the transition. The maps associated to the other cities are shown in figs. S37-S44.

C. Path performance inequality

Studying path performance from a socio-economical perspective, we could ask ourselves how the network degradation impacts the whole population of paths as congestion grows. It is interesting to study whether paths are uniformly affected or there is an imbalance leading to some paths still being traversable efficiently while the majority is catastrophically broken. Of course, this question makes particular sense near the transition where we know that the system's susceptibility to traffic changes is maximum. In fig. 10 (see figs. S45-S46 for the other cities) we see some specific P distributions (four vertical slices from fig. 8(bottom)) where we can better observe how P shifts from an average high performance (solid line) at 50% of the transition that flattens at the flex and immediately beyond (dot-dashed and dashed lines). In fig. 11 ((see figs. S47-S48 for the other cities)), we plot the Gini coefficient at all traffic levels (vertical lines are placed at the same reference levels as in fig. 10): at half the flex level the degradation is still moderate and uniform (black solid line) while at the flex (dash-dotted) the distribution flattens and spans the whole P range and the Gini coefficient grows to ~ 0.4 . For even higher traffic we see a P distribution peaked on extremely high path degradation and a Gini over 0.5, meaning that very few paths are still able to perform well and the rest are almost unusable. The Gini coefficient acts also as a good proxy for detecting the transition because it is much more sensitive to performance degradation than the path rejection ratio that we used in fig. 5: The flex position is detected at the same traffic value, though.

IV. CONCLUSION

We have investigated how the fastest paths on transportation networks evolve as congestion levels increase, using a proven model in which speed decreases linearly with density at the edge level due to interactions among selfish agents. In a different setting, this model has shown the ability to grasp important structural properties of real urban traffic [9].

In this work, we significantly extend the analysis, shifting the measuring process from an edge-based to a path-based network performance metric. Studying urban vehicular transportation with the selected model, is particularly important for city analysis since the emergence of personal navigation tools, providing the fastest routes for their users in real-time, has shown to have a major impact on traffic patterns [10]. We cover both random planar networks with controlled characteristics and




FIG. 9. P maps for London: Top (bottom) plots show the average P value associated with the Origin (Destination) nodes of each path.




FIG. 10. London P distributions at four traffic levels, three of which can be directly compared to the P maps in fig. 9.

real urban networks of major cities, and show the emergence of complex and spatially non-uniform degradation of network performance under growing traffic. In order to quantify this degradation, we introduced a novel measure called Performance Index, which incorporates path slowdown and completeness.

We started by analyzing the low-density network regime, in which an initial smooth path performance degradation is observed up to a critical traffic volume, and is followed, with higher traffic, by the emergence of disjoint subnetworks: This catastrophic decrease of network performance turns out to be due to the failure of a




FIG. 11. The Gini coefficient as traffic load increases for the city of London. The same traffic levels of fig. 9 with the same line styles are plot as vertical lines for comparison.

very small number of edges.

By characterizing the effective length, detour, and inness for the fastest paths, as well as analyzing the performance degradation distribution across uniformly distributed paths, we have shown the interplay between traffic growth and network efficiency. Fastest paths for all cities show a peak for detour and inness (and their variance) in the proximity of the critical traffic level. We locate the transition to congestion at the flex of the hard-rejected path ratio curve. For most urban areas, inness uncovers, with light traffic, that paths are slightly attracted by city centers. On the other hand, this central force becomes strongly repulsive during congestion.

By exploiting the proposed path performance index, we demonstrated an important asymmetry in the efficiency of some parts of the cities when operating as traffic sources or as sinks. Finally, the Gini coefficient, applied to path performance, was established as a useful precursor for the congestion transition and also as a simple measure to quantify the inequality of degradation over different parts of the network.

Our results showed that, in order to build resilient cities to better exploit their maximum traffic volume potential, naturally involves having multiple available paths among adjacent neighborhoods (that are usually well connected) avoiding single points of failure that tend to saturate early regardless of their capacity. Several smaller connections are usually more resistant than a

single one with equivalent capacity because spatial multiplicity tends to distribute traffic over larger areas, better exploiting the available potential road volume.

These results promise to offer valuable insights into the vulnerabilities of transportation networks under increasing congestion. Note that, even though in this work our method was used to analyze patterns typical of urban vehicular traffic, it is expected that other transport phenomena involving agent competition for network resources could be approached similarly.

Acknowledgments – We are deeply indebted to Francesco Versaci for the many useful discussions. We acknowledge the contribution of Sardinian Regional Authorities under project XDATA (art 9 L.R. 20/2015). Map data copyrighted OpenStreetMap contributors: www.openstreetmap.org.

-
- [1] Dirk Helbing. Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.*, 73(4):1067–1141, December 2001. Publisher: American Physical Society.
 - [2] Guanwen Zeng, Daqing Li, Shengmin Guo, Liang Gao, Ziyu Gao, H. Eugene Stanley, and Shlomo Havlin. Switch between critical percolation modes in city traffic dynamics. *Proc Natl Acad Sci USA*, 116(1):23–28, January 2019.
 - [3] Mengyao Zhang, Tao Huang, Zhaoxia Guo, and Zhenggang He. Complex-network-based traffic network analysis and dynamics: A comprehensive review. *Physica A: Statistical Mechanics and its Applications*, 607:128063, 2022.
 - [4] Clémence Magnien, Matthieu Latapy, and Jean-Loup Guillaume. Impact of random failures and attacks on Poisson and power-law random networks. *ACM Comput. Surv.*, 43(3), April 2011.
 - [5] Milena Oehlers and Benjamin Fabian. Graph metrics for network robustness—a survey. *Mathematics*, 9(8):895, 2021.
 - [6] Alice C Schwarze, Jessica Jiang, Jonny Wray, and Mason A Porter. Structural robustness and vulnerability of networks. *arXiv preprint arXiv:2409.07498*, 2024.
 - [7] Guanwen Zeng, Jianxi Gao, Louis Shekhtman, Shengmin Guo, Weifeng Lv, Jianjun Wu, Hao Liu, Orr Levy, Daqing Li, Ziyu Gao, H. Eugene Stanley, and Shlomo Havlin. Multiple metastable network states in urban traffic. *Proc Natl Acad Sci USA*, 117(30):17528–17534, July 2020.
 - [8] Marco Cogoni and Giovanni Busonera. Stability of traffic breakup patterns in urban networks. *Physical Review E*, 104(1):L012301, 2021.
 - [9] Marco Cogoni and Giovanni Busonera. Predicting network congestion by extending betweenness centrality to interacting agents. *Physical Review E*, 109(4):L044302, 2024.
 - [10] Nimrod Serok, Orr Levy, Shlomo Havlin, and Efrat Blumenfeld-Lieberthal. Unveiling the inter-relations between the urban streets network and its dynamic traffic flows: Planning implication. *Environment and Planning B: Urban Analytics and City Science*, 46(7):1362–1376, September 2019. Publisher: SAGE Publications Ltd STM.
 - [11] Petter Holme. Congestion and centrality in traffic flow on complex networks. *Advances in Complex Systems*, 6(02):163–176, 2003.
 - [12] Roger Guimera, Stefano Mossa, Adrian Turttschi, and LA Nunes Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799, 2005.
 - [13] Aisan Kazerani and Stephan Winter. Can betweenness centrality explain traffic flow. In *12th AGILE international conference on geographic information science*, pages 1–9, 2009.
 - [14] Chi Ho Yeung and David Saad. Competition for shortest paths on sparse graphs. *Physical review letters*, 108(20):208701, 2012.
 - [15] Homayoun Hamedmoghadam, Mahdi Jalili, Hai L. Vu, and Lewi Stone. Percolation of heterogeneous flows uncovers the bottlenecks of infrastructure networks. *Nat Commun*, 12(1):1254, February 2021.
 - [16] Luis E. Olmos, Serdar Çolak, Sajjad Shafiei, Meead Saberi, and Marta C. González. Macroscopic dynamics and the collapse of urban traffic. *Proceedings of the National Academy of Sciences*, 115(50):12654–12661, December 2018. Publisher: Proceedings of the National Academy of Sciences.
 - [17] Serdar Çolak, Antonio Lima, and Marta C. González. Understanding congested travel in urban areas. *Nat Commun*, 7(1):10793, April 2016.
 - [18] Alexandre Diet and Marc Barthelemy. Towards a classification of planar maps. *Phys. Rev. E*, 98(6):062304, December 2018. Publisher: American Physical Society.
 - [19] David Aldous and Karthik Ganesan. True scale-invariant random spatial networks. *Proceedings of the National Academy of Sciences*, 110(22):8782–8785, 2013.
 - [20] Daqing Li, Bowen Fu, Yunpeng Wang, Guangquan Lu, Yehiel Berezin, H. Eugene Stanley, and Shlomo Havlin. Percolation transition in dynamical traffic network with evolving critical bottlenecks. *Proc Natl Acad Sci USA*, 112(3):669–672, January 2015.
 - [21] Alec Kirkley, Hugo Barbosa, Marc Barthelemy, and Gourab Ghoshal. From the betweenness centrality in

- street networks to structural invariants in random planar graphs. *Nat Commun*, 9(1):2501, June 2018. Number: 1 Publisher: Nature Publishing Group.
- [22] Minjin Lee, Hugo Barbosa, Hyejin Youn, Petter Holme, and Gourab Ghoshal. Morphology of travel routes and the organization of cities. *Nat Commun*, 8(1):2229, December 2017. Number: 1 Publisher: Nature Publishing Group.
- [23] Hesham Rakha and Brent Crowther. Comparison of greenshields, pipes, and van aerde car-following and traffic stream models. *Transportation Research Record*, 1802(1):248–262, 2002.
- [24] Rui Ding, Norsidah Ujang, Hussain bin Hamid, Mohd Shahrudin Abd Manan, Yuou He, Rong Li, and Jianjun Wu. Detecting the urban traffic network structure dynamics through the growth and analysis of multi-layer networks. *Physica A: Statistical Mechanics and its Applications*, 503:800–817, 2018.
- [25] Alexander P Kartun-Giles, Marc Barthelemy, and Carl P Dettmann. Shape of shortest paths in random spatial networks. *Physical Review E*, 100(3):032315, 2019.
- [26] Marco Cogoni, Giovanni Busonera, and Gianluigi Zanetti. Ultrametricity of optimal transport substates for multiple interacting paths over a square lattice network. *Physical Review E*, 95(3):030108, 2017.