



Cite this article: Williams MJ, Musolesi M.
2016 Spatio-temporal networks: reachability,
centrality and robustness. *R. Soc. open sci.*
3: 160196.
<http://dx.doi.org/10.1098/rsos.160196>

Received: 17 March 2016

Accepted: 24 May 2016

Subject Category:

Computer science

Subject Areas:

computer modelling and
simulation/complexity/applied mathematics

Keywords:

data science, complex systems, urban
analytics, transport systems, network
robustness, urban data science

Author for correspondence:

Matthew J. Williams
e-mail: m.j.williams@cs.bham.ac.uk

Electronic supplementary material is available
at <http://dx.doi.org/10.1098/rsos.160196> or via
<http://rsos.royalsocietypublishing.org>.

Spatio-temporal networks: reachability, centrality and robustness

Matthew J. Williams^{1,2} and Mirco Musolesi^{2,1}

¹School of Computer Science, University of Birmingham, Edgbaston B15 2TT, UK

²Department of Geography, University College London, London WC1E 6BT, UK

 MJW, 0000-0003-0892-0998

Recent advances in spatial and temporal networks have enabled researchers to more-accurately describe many real-world systems such as urban transport networks. In this paper, we study the response of real-world spatio-temporal networks to random error and systematic attack, taking a unified view of their spatial and temporal performance. We propose a model of spatio-temporal paths in time-varying spatially embedded networks which captures the property that, as in many real-world systems, interaction between nodes is non-instantaneous and governed by the space in which they are embedded. Through numerical experiments on three real-world urban transport systems, we study the effect of node failure on a network's topological, temporal and spatial structure. We also demonstrate the broader applicability of this framework to three other classes of network. To identify weaknesses specific to the behaviour of a spatio-temporal system, we introduce centrality measures that evaluate the importance of a node as a structural bridge and its role in supporting spatio-temporally efficient flows through the network. This exposes the complex nature of fragility in a spatio-temporal system, showing that there is a variety of failure modes when a network is subject to systematic attacks.

1. Introduction

Network science provides many powerful methods to study a great variety of systems in society, nature and technology. Modelling complex systems in terms of their network structure allows researchers to understand, predict and optimize their real-world behaviour [1,2]. Detailed data describing the interactions and relationships in real-world systems have become increasingly available, emerging from domains as diverse as transport [3], biology [4,5], infrastructure [6,7] and sociology [8]. The ability of network analysis to capture relationships and dependencies between components makes it an essential tool for studying system resilience [9–13]. It is known that local disruptions

can have a significant impact on the global behaviour of a system [9,14]. The study of a network's vulnerabilities and its response to malfunction helps engineers design robust systems [7,14–18] and scientists understand complex phenomena such as neurodysfunction [19,20], economic and financial risk [21,22], and disease spreading [23].

Methods for robustness analysis generally assume that the system is represented by a network composed of static links, focusing on the topological properties of a network subject to disruption. Depending on the system and research question, a static representation may also incorporate weighted and directed edges, allowing richer dynamics to be modelled. In many systems, however, edges are not continuously active [24] and the quantities their weights represent may vary with time. Furthermore, these time-varying systems may also be spatially embedded, and thus the ability for nodes to interact is governed by the space in which they operate as well as their network connectivity.

Urban transport systems have a long history in network analysis [11,25–29] and are an exemplar of a system that is temporal [24], spatial [30] and multilayer [31]. Examples of recent studies of transport systems that account for time-varying and multilayer properties, either jointly or separately, can be found in [11,25,29,32]. Furthermore, the resilience of such systems is of particular interest, with insights from static network analysis enabling better-engineering of, for example, rail and metro systems [28,33]. Aside from transport systems, there are many classes of network that are spatial, temporal or both. Engineering spatial network resilience in fixed communication networks has particularly received attention over recent years, including new methods for identifying critical geographical regions [34–36] and understanding spatial damage to fibre-optic networks [37]. Communication networks and cellular nervous systems are just few examples of systems that can be represented by means of networks embedded in space and time. Relying only on a static space-agnostic aggregation of such networks over-simplifies the rich and complex relationships in the real systems they represent.

The consequences of ignoring the temporal and spatial constraints on networks have been highlighted in recent surveys of spatial [30], temporal [24] and multilayer [31,38] networks. Advances in temporal networks have provided researchers with a valuable framework that can be used to understand the temporal structure of systems, and empirical measurements have demonstrated that static representations often over-estimate the true connectivity of real-world temporal systems [39–42]. Furthermore, while static methods investigate purely topological measures of system vulnerability, such as giant component size and network diameter, new methods allow researchers to additionally explore temporal efficiency [43] and expose differences in the temporal function of the network when subject to random error (Err) versus systematic attack [40,44].

However, temporal network models sometimes make the strong assumption that interactions are instantaneous. For many types of spatially embedded system, this assumption ignores the influence that space has in constraining the structure of the network. As an example, transit between stations in a public transport system naturally incurs a time delay while a passenger travels, and the specific delay depends on the speed of the service currently operating and distance between stations. In one of the earliest works on temporal networks, Kempe *et al.* [45] propose a schedule-based analysis of transport systems as a potential application of temporal network analysis; more-recent application of time-respecting paths [45,46] to transport networks can be found in Pan & Saramäki [32], and Gallotti & Barthelemy [25]. Critical nodes in such a system are those that not only bridge two clusters in the network, but also act as a conduit for *rapid* flow between physically distant areas of the system.

To study robustness in spatially embedded temporal systems, we apply a general framework that is able to capture instantaneous and non-instantaneous types of interaction. We formulate a model of spatio-temporal systems in which the interactions and relationships between components are constrained by the space and time in which they are embedded (§2). The analysis of a network's spatio-temporal structure (defined in §2) provides the foundation to measure its topological, temporal and spatial function. These measures are presented in §3, along with systematic attack strategies designed to expose different weaknesses in the network. In order to validate our framework and show its utility and flexibility, in this paper we explore the behaviour and resilience of six empirical examples of spatio-temporal systems (described in §4); specifically, three urban transit systems, a national air travel network, a biological neural network and a mobile phone communication network. Our empirical analysis in §5 highlights a crucial distinction when examining the robustness of spatio-temporal systems: we must consider the impact of disruption to a network's temporal and spatial efficiency, in addition to its topological structure. Furthermore, critical nodes play different roles in terms of their topological, temporal and spatial utility, and therefore systematic attack strategies can differ in the damage they cause to the network.

2. Connectivity in spatio-temporal networks

We start by introducing the notation and framework with which we define connectivity in a spatio-temporal network. Our model follows a fundamental property of spatial networks: the space in which the system is embedded acts as a constraint on the structure of the network [30]. In a spatio-temporal setting, we can represent this constraint as the speed with which one node can interact with another. This is a natural abstraction for many real-world processes that have previously been modelled as purely static or temporal systems; for example, passenger transit in public transport systems, neurotransmission in biological neural networks (e.g. [47]), shipping in multimodal freight networks (e.g. [48]), movements in trade networks (e.g. [49]) and signal delay in telecommunications networks (such as the Internet). Our approach mixes structure (i.e. topology), space and time, and allows for an exploration of the influence of each dimension on the processes occurring in these networks.

We summarize the relevant notation introduced throughout §§2–3 and in table 1.

2.1. Time-varying spatially embedded systems

A time-varying network is conventionally represented as a time-ordered sequence of graphs [24,45,50–52], with each graph corresponding to a snapshot of the network during a particular time window. The time intervals are commonly finite and equally sized, and we refer to the interval duration as the *temporal granularity* τ . We assume an overall observation duration that consists of T timesteps, starting with an initial time t_1 .

More formally, let V denote the set of all nodes in the system under study, and let $N = |V|$ be the number of nodes in the system. We consider a temporal graph consisting of T discrete non-overlapping windows, represented by the time-ordered sequence of directed graphs $G^{[t_1]}, \dots, G^{[t_T]}$. A particular graph $G^{[t]} = (V, E^{[t]})$ captures the topological state of the spatio-temporal system during the interval $[t, t + \tau)$. It is assumed that nodes are present throughout the lifetime of the system; that is, each graph has the same node set V . For each graph $G^{[t]}$, where $t = t_1, \dots, t_T$, there exists a counterpart weight matrix $\mathbf{S}^{[t]} \in \mathbb{R}^{N \times N}$ that represents the weighted directed edges between nodes during the interval $[t, t + \tau)$. The weight of the edge from node v to node w at time t corresponds to the element $\mathbf{S}_{vw}^{[t]}$. In our model, $\mathbf{S}_{vw}^{[t]}$ is a non-negative scalar representing the speed of physical propagation from v to w in the corresponding time interval. For example, in a public transport system, v and w are typically transit stations and $\mathbf{S}_{vw}^{[t]}$ is the average speed at which passengers are conveyed from v to w given the service operating during $[t, t + \tau)$. In the case that there is no connection from node v to w at time t , we have $\mathbf{S}_{vw}^{[t]} = 0$.

In a spatio-temporal system, the constituent entities naturally occupy a location in space, and this location may vary over time. More formally, we assume nodes are embedded in a k -dimensional metric space with physical distance function g . The embedding is dependent on the system; for example, geographical networks are commonly represented in a spherical (or ellipsoidal) coordinate system (where g is geodesic distance over the surface of the Earth) or projected on to a two-dimensional Euclidian space. We write $l_v^{[t]}$ to denote the physical position¹ of node v during the interval $[t, t + \tau)$, noting that $l_v^{[t]} \in \mathbb{R}^k$. For convenience, we collect the time-varying pairwise distances between nodes in a physical distance matrix $\mathbf{D}^{[t]}$, where $\mathbf{D}_{vw}^{[t]} = g(l_v^{[t]}, l_w^{[t]})$ for each $v, w \in V$. We make the simplifying assumption that multiple nodes in the network cannot occupy the same location. In real-world systems, nodes represent physical entities, and thus cannot occupy the same space.

2.2. Model description

To apply classical network-theoretic concepts in a spatio-temporal system, we use the notion of space-time constrained propagation between nodes. Paths through the spatio-temporal network are defined by this process, and then form the basis for higher-order network measures such as connected components and network distance. These paths obey the time-varying conditions of the system along their route. This form of connectivity follows the same spreading process that is common in defining temporal paths [50–52], with the modification that propagation from one node to another is constrained by the speed of transmission between nodes and their physical distance. This framework's explicit encoding of time-varying speeds serves as an alternative to event-based representations of similar networks (e.g. as applied to an air transport network in [32]).

¹In practice, a non-stationary node may occupy many locations during this interval, and the method for choosing a single representative location $l_v^{[t]}$ depends on the system being studied.

Table 1. Summary of notation used in this paper.

symbol	description
τ	window width
$G^{[t]}$	graph for snapshot at time t
$\mathbf{S}^{[t]}$	matrix of physical propagation speeds between nodes for snapshot at time t
$\mathbf{D}^{[t]}$	matrix of physical distances between nodes for snapshot at time t
$K^{[t]}$	reachability set (dependent on implicit origin node v_0 and start time t_1)
$\mathbf{P}^{[t]}$	progress matrix (dependent on implicit origin node v_0 and start time t_1)
λ	temporal efficiency
ζ	spatial efficiency
S	size of largest strongly connected component as fraction of all nodes
$R^\lambda(f)$	temporal robustness at deactivation rate f
$R^\zeta(f)$	spatial robustness at deactivation rate f

We consider propagation as a discrete-time process starting at an origin node v_0 and the initial time t_1 , and progressing over each subsequent timestep t_2, t_3, \dots, t_T . In the following description of the propagation process, we treat the origin node v_0 and start time t_1 as implicit parameters. Modelling non-instantaneous propagation in a time-varying system necessarily involves capturing partial propagation between nodes, as well as the nodes that have been reached along a path. The process is represented by two time-evolving structures: a *reachability set* $K^{[t]}$ and an N -by- N *progress matrix* $\mathbf{P}^{[t]}$. The reachability set $K^{[t]}$ consists of nodes that have been reached from origin node v_0 by the end of timestep t . The progress matrix $\mathbf{P}^{[t]}$ represents the distance accumulated during direct propagation between nodes; more specifically, the element $\mathbf{P}_{vw}^{[t]}$ gives the *progress* from node v to node w at timestep t measured as the distance accumulated during propagation from v to w . A progress element represents the state of partial propagation between the two nodes at the end of a timestep. In this process, progress from node v to w is able to accumulate while the two nodes are continuously connected by an edge. The amount by which progress is incremented in each timestep is governed by the propagation speed between the two nodes. Given a sufficient period of continuous connectivity between v and w , progress can accumulate to the point where it exceeds the physical distance between the two nodes, thereby representing complete propagation from v to w .

To give an example, consider the case of a national rail network where the elements of the network are cities and edges are transport links. Let us consider transportation with vehicles travelling between two cities v and w at an average speed of 200 km h^{-1} . If the distance between the cities is 150 km , after a 15 min time interval (indicated by $\tau = 15 \text{ min}$) the value $\mathbf{P}_{vw}^{[t]}$ will be equal to 50 km . Thus, if the average speed remains constant, then three timesteps must elapse before w is reachable from v .

Let us now consider the more-complex example in figure 1. Here we can see that source node A reaches B and C in at most three subsequent timesteps. Specifically, full propagation from A to B is able to occur in one timestep, arriving at timestep t_2 . Then, in the case of propagation from B to C , the distance between the nodes is larger, and thus more time is required to complete the interaction. We therefore see that there is an intermediate stage in timestep t_3 where there is only partial propagation from B to C . This completes in timestep t_4 . The example also demonstrates two failed propagation attempts between A and D . In t_3 propagation is set to zero due to the absence of an edge from A to D . The reappearance of the edge in timestep t_4 allows propagation to restart, but subsequently fails again at t_6 .

We now formalize the propagation process. The initial state of the process is such that origin node v_0 is the only reachable node and no progress has been made; that is, at the initial timestep t_1 we have $K^{[t_1]} = \{v_0\}$ and $\mathbf{P}^{[t_1]}$ is all zero. When progressing from timestep t_{i-1} to t_i , where $i \geq 2$, further propagation from a reachable node v to an unreachable node w depends on whether an edge from v to w exists in timestep t_i . If these conditions hold, the distance by which we can increment progress from v to w is the *propagation increment* $\tau \cdot \mathbf{S}_{vw}^{[t_i]}$, where $\mathbf{S}_{vw}^{[t_i]}$ is propagation speed and τ is the window duration. We note that the propagation increment may exceed the *remaining distance* required to propagate from v to w , which we denote by $q_{vw}^{[t_i]}$. The quantity $\tau \cdot \mathbf{S}_{vw}^{[t_i]}$ therefore represents the maximum amount that may be added to

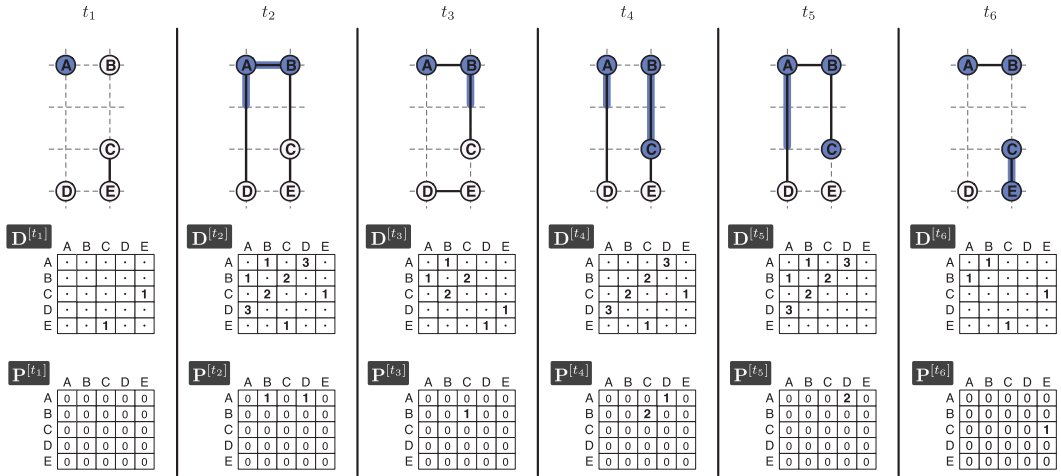


Figure 1. Example of the constrained propagation process in a spatio-temporal system. Propagation begins with origin node A at timestep t_1 . Nodes are positioned at integer locations on a Euclidian plane. Temporal granularity $\tau = 1$ s. All edges have a propagation speed of 1 m s^{-1} . Membership of a node in the reachability set $K^{[t]}$ at the end of timestep t is represented by shading. The corresponding distance matrix $\mathbf{D}^{[t]}$ and progress matrix $\mathbf{P}^{[t]}$ are shown below each network snapshot. Progress is also visually represented by shading along an edge. For clarity, we omit from $\mathbf{D}^{[t]}$ distances between pairs of nodes that are not connected in $G^{[t]}$.

the progress of v to w , while the actual increment is given by

$$\min(\tau \cdot \mathbf{S}_{vw}^{[t_i]}, q_{vw}^{[t_i]}). \tag{2.1}$$

The remaining distance $q_{vw}^{[t_i]}$ is obtained from the current physical separation distance $\mathbf{D}_{vw}^{[t_i]}$ and previous progress $\mathbf{P}_{vw}^{[t_i-1]}$ as

$$q_{vw}^{[t_i]} = \mathbf{D}_{vw}^{[t_i]} - \mathbf{P}_{vw}^{[j-1]}, \tag{2.2}$$

or set to 0 if this quantity becomes negative due to a change in the pair's positions between the two timesteps.

Finally, we derive the progress matrix update rule. The value $\mathbf{P}_{vw}^{[t_i]}$ in timestep t_i , where $i \geq 2$, is

$$\mathbf{P}_{vw}^{[t_i]} = \mathbf{P}_{vw}^{[t_i-1]} + \min(\tau \cdot \mathbf{S}_{vw}^{[t_i]}, q_{vw}^{[t_i]}) \tag{2.3}$$

if $v \in K^{[t_i-1]}$, $w \notin K^{[t_i-1]}$, and edge (v, w) exists in $G^{[t_i]}$. In the other case that the system has no edge from v to w at timestep t_i , progress is reset and so $\mathbf{P}_{vw}^{[t_i]}$ is set to 0. Propagation from v to w is considered successful in timestep t_i if the cumulative progress $\mathbf{P}_{vw}^{[t_i]}$ between the nodes exceeds or equals their physical distance $\mathbf{D}_{vw}^{[t_i]}$. In such a case, w is regarded as reachable from source node v_0 at timestep t_i , and w is included in the set $K^{[t_i]}$. More formally, the reachability set at timestep t_i is expressed as

$$K^{[t_i]} = K^{[t_i-1]} \cup \{w \mid \exists v \text{ s.t. } \mathbf{P}_{vw}^{[t_i]} \geq \mathbf{D}_{vw}^{[t_i]}\}. \tag{2.4}$$

In the preceding equation, $\mathbf{P}_{vw}^{[t_i]} \geq \mathbf{D}_{vw}^{[t_i]}$ represents the case that sufficient time has elapsed for propagation to complete, expressed in terms of the cumulative propagation distance $\mathbf{P}_{vw}^{[t_i]}$ with respect to the current separation $\mathbf{D}_{vw}^{[t_i]}$ between the two nodes. Hence, the process models the effect of nodes that may be non-stationary in their spatial embedding. This means that, for example, propagation is accelerated when two nodes are moving towards one another.

To summarize, the spatio-temporal constrained propagation process is entirely described by the time-evolution of three matrices: a physical distance matrix $\mathbf{D}^{[t]}$ of spatial distances between nodes; a propagation speed matrix $\mathbf{S}^{[t]}$, which specifies the speed of transmission between pairs of connected nodes; and a progress matrix $\mathbf{P}^{[t]}$, which tracks incremental propagation between nodes and evolves according to equation (2.3). Node-to-node propagation is successful only after sufficient time has elapsed, and this duration is a result of the timestep width τ , the time-dependent distance between nodes and the time-dependent speed between nodes.

As with discrete-time temporal network representations (e.g. [47,53–56]), the temporal granularity τ represents a necessary trade-off between abstraction and fidelity. In our propagation process, finer granularities minimize the amount of excess in the propagation increment on successful propagation

(i.e. $\mathbf{P}_{vw}^{[t_i]} - \mathbf{D}_{vw}^{[t_i]}$). Too-coarse a granularity leads to under-estimation of the true propagation duration, and over-estimation of the propagation distance. An upper bound for the quality of the choice of the parameter τ can be obtained by considering the physical distances and speeds of the underlying system. In particular, we find the minimum non-zero value of the product $\mathbf{D}_{vw}^{[t]} \cdot \mathbf{S}_{vw}^{[t]}$ over all $t = t_1, \dots, t_T$ and $v, w \in V$. For the given choice of τ , this quantity represents the minimum direct propagation duration that exists in any of the system's snapshots. If the minimum direct propagation time is greater than τ , it means that the representation is guaranteed to under-sample the temporal dynamics of the network.

We have so far considered propagation through the system from a single origin node v_0 , starting at timestep t_1 , and proceeding through timesteps $t = t_2, t_3, \dots, t_T$. For network analysis, this allows us to construct spatio-temporal paths originating at v_0 . The main features from the propagation process we consider are the time at which a node w was reached from v_0 , and the physical distance travelled through intermediate propagation along the route from v_0 to w . By repeating this process from each node in V , we construct system-wide network connectivity in the time interval $[t_1, t_T + \tau)$, thereby yielding spatio-temporal paths starting at time t_1 .

2.3. Spatio-temporal paths and distance measures

In general, a *spatio-temporal path* from node v_0 may visit multiple distinct vertices before reaching its destination node. A spatio-temporal path consisting of $n \geq 0$ hops, starting with origin node v_0 at timestep t_1 , is described as the sequence of $n + 1$ pairs

$$((v_0, t_1), (v_1, t'_{\text{arr}_1}), (v_2, t'_{\text{arr}_2}), \dots, (v_n, t'_{\text{arr}_n})), \quad (2.5)$$

where v_j denotes the j th node visited on the path and t'_{arr_j} denotes the time at which the path reached node v_j . Spatio-temporal paths are readily constructed by tracing the propagation process described in the previous section.

In addition to the timestep t'_{arr_j} in which the path arrives at a node v_j , the propagation process also captures the physical distance traversed in reaching v_j from preceding node v_{j-1} , given by the progress matrix element

$$\mathbf{P}_{v_{j-1}v_j}^{[t'_{\text{arr}_j}]} \quad (2.6)$$

This quantity is useful for studying the distance travelled along the route. We also note that any such path constructed from the propagation process also obeys the temporal-ordering condition

$$t_1 < t'_{\text{arr}_1} < \dots < t'_{\text{arr}_n}. \quad (2.7)$$

Furthermore, as is the case with temporal networks, reachability in spatio-temporal networks is non-transitive and non-symmetric.

To illustrate an example, in [figure 1](#) we can see that node A is able to reach nodes B , C and E . From this, we can construct the spatio-temporal routes taken to reach each node from A . For example, to reach E we can trace the path from A through nodes B and C and identify the total time taken as 5 s and overall physical distance travelled along the path as 4 m.

Throughout the rest of the paper, our notation will omit the observation window $[t_1, t_T + \tau)$ in which the spatio-temporal paths are obtained from a spatio-temporal system, leaving it as an implicit parameter. We stress, however, that the following measures we introduce are defined on the paths present in such an interval.

Temporal paths are characterized by two notions of length [24,50]: the *topological length*, which is the number of hops along the path,² and *latency*, which is the time elapsed from source to destination. In reference to the previous description of a spatio-temporal path (equation (2.5)), topological length corresponds to n and latency is given by $t'_{\text{arr}_n} - t_1$. In the context of a spatio-temporal path, an additional relevant feature is the *spatial length* of a path from v_0 to v_n , given by

$$\sum_{j=1}^n \mathbf{P}_{v_{j-1}v_j}^{[t'_{\text{arr}_j}]} \quad (2.8)$$

These quantities allow us to explore connectivity in the network from topological, temporal and spatial perspectives.

²Specifically, topological length is the number of hops along the (temporally) shortest path, which may not necessarily be a path with the minimum number of hops between the two nodes.

To study the robustness of the system, we focus on the shortest paths between nodes, as these represent optimal routes within the system. Indeed, the typical length of such paths is assumed to be representative of the overall efficiency of the network. As commonly defined in temporal networks [24,50], a path from node v to w is a temporally shortest path if it has minimum latency. We extend this definition to our spatio-temporal setting. Formally, a path from v to w is a *spatio-temporally shortest path* if it is a temporally shortest path and has minimum spatial length.

Although multiple temporally shortest paths may exist from one node to another, it is a subset of these paths that are spatio-temporally shortest, and all such paths share the same latency and spatial length. Under this definition, notions of spatio-temporal distance in the network have a well-defined meaning. Specifically, the *temporal distance* dist_{vw}^λ from v to w is the latency of the spatio-temporal shortest path from v to w . Similarly, the *spatial distance* dist_{vw}^s from v to w is the spatial length of the spatio-temporal shortest path from v to w . In the case that w cannot be reached from v in the given time interval (and thus no path exists between them) both distance quantities are set to ∞ .

In practice, to compute spatio-temporal shortest paths in a network, we follow an epidemic (i.e. breadth-first) approach, similar to the Brandes method [57] for computing betweenness centrality in weighted static networks. As with the Brandes algorithm, we must enumerate shortest paths through the network; however, due to the non-transitive and non-symmetric nature of temporal and spatio-temporal paths, intermediate shortest paths cannot be re-used in the calculation of broader paths. Here we construct a breadth-first shortest path tree for each node in the network. The tree is rooted at the given node, and spatio-temporal shortest paths from the root are obtained according to the process defined in §2.2. In our implementation, computing each tree can require a worst-case time of $\mathcal{O}(N^2)$ where N is the number of nodes. In our practical experiments (see datasets in §4), we parallelize this tree-construction process, and find that it is feasible to extract shortest paths in networks up to a size of 400–450 nodes, though note that runtime also depends on the sparsity of the network.

3. Random error and systematic attack

In order to study the resilience of a network we define different ways in which nodes are selected to fail.

Err refers to probabilistic models of node failure. In this paper, we consider uniform random failure, where each node has an equal independent failure probability of f , yielding an expected number of node failures $f \cdot N$. *Err* allows us to measure the typical response of the network to random breakdowns. It is also a useful comparison for systematic attack, as any intelligent attack strategy should be at least as harmful to the network as *Err*.

Systematic attacks on networks rely on the ability to strategically identify and deactivate nodes that are critical to the function of the system. Although many measures of node centrality are available to define attack strategies [1], only a subset of these are effective in rapidly disrupting the behaviour of a network [32,40]. Attack on *spatio-temporal* networks adds further complexity, as the importance of a node can be evaluated in terms of its centrality in the topological, temporal and/or spatial structure of the system, properties which are not necessarily correlated.

In the rest of this section, we introduce relevant centrality measures and formalize the attack strategies we use to test the robustness of a spatio-temporal network. We also define the quantities with which we measure the performance of a network.

3.1. Spatio-temporal centrality measures

A variety of formulations of closeness in temporal networks have been developed. Given the non-symmetry of connectivity and latency in temporal networks, we must distinguish between in- and out-closeness. In particular, common definitions include temporal out-closeness [32] and temporal in-farness [40,44]. The latter has been proposed as a basis for a centrality measure in systematic attacks on networks [40,44]. Here we present temporal in-closeness as the reciprocal of temporal in-farness. Formally, given the giant temporal in-component \mathcal{C}_v of node v , we can consider the average temporal distance of nodes that reach v as

$$\ell_v = \frac{1}{|\mathcal{C}_v| - 1} \sum_{w \in \mathcal{C}_v; w \neq v} \text{dist}_{vw}^\lambda. \quad (3.1)$$

A node with low ℓ_v is considered more-central in the network. An unreachable node has no utility in the network, and thus we set $\ell_v = \infty$ if \mathcal{C}_v is singleton. The *temporal in-closeness* of a node v is then




Figure 2. Comparison of path betweenness (PB) and betweenness efficiency (BE). Although both shaded nodes have the same PB, the bottom node has higher BE. Deactivating the bottom node has a larger impact on the overall temporal performance of the network.

given by

$$C_v^{\text{TC}} = \frac{1}{\ell_v}. \quad (3.2)$$

Inverting in-farness simply reverses the ranking of nodes, ensuring high values for more-central nodes.

However, while closeness is useful for identifying nodes capable of rapidly reaching the rest of the network (e.g. for fast patching of computer malware [58]), it does not directly measure the reliance of the network on a node as an efficient connector, which is often a critical source of network fragility. This notion is better represented by betweenness centrality [12], which we extend to our spatio-temporal network as follows. First, we denote the set of spatio-temporally shortest paths from node w to node u by σ_{wu} and let $\sigma_{wu}(v)$ denote the subset of paths in σ_{wu} that pass through node v . *Path betweenness (PB) centrality* is then defined for a node v as

$$C_v^{\text{PB}} = \sum_{w,u \in V; w \neq u} \frac{|\sigma_{wu}(v)|}{|\sigma_{wu}|}. \quad (3.3)$$

Although this quantity represents the importance of a node in supporting connectedness within the network, it does not necessarily reflect a node's potential role as a hub through which information can be quickly transferred. Indeed, efficient temporal conduits such as these are especially important for rapid transmission within a network (as recently highlighted in [55]), and deactivating these conduits can have significant impact on the global efficiency of the network. We therefore define the *betweenness efficiency (BE) centrality* of a node v as

$$C_v^{\text{BE}} = \sum_{(w,u) \in \beta_v} \frac{1}{\text{dist}_{wu}^\lambda}, \quad (3.4)$$

where β_v is the set of source–destination pairs for each path in $\sigma_{wu}(v)$. The reciprocal distance $1/\text{dist}_{wu}^\lambda$ is useful as a measure of the efficiency of the shortest path w to u , and is commonly used in evaluating the network efficiency of static [59,60] and temporal [24,50,61] networks.

3.2. Attack strategies

We summarize our five chosen attack strategies as follows. Temporal closeness (TC) attack deactivates nodes in decreasing order by temporal in-closeness centrality. This follows the order defined in the initial (intact) network, similar to the approach in [40].

PB and BE attacks exploit PB and BE centralities, respectively. As shown in [12], the geodesic paths in a network can significantly change after each failure. Thus, in both betweenness-based attacks, the centrality ranking is recalculated after each node failure.

Figure 2 demonstrates the difference in how PB and BE attacks prioritize their node deactivations. Deactivating either shaded node will have a similar effect on the network's spatio-temporal connectivity; however, the bottom node is supporting paths that have faster propagation flow. BE attack prioritizes nodes which are bottlenecks with high temporal load.

Finally, we include two attacks based on the degree distribution of the intact network: in-degree (ID) and out-degree (OD) attacks. Unlike the other three attack strategies, the ID and OD strategies do not rely on global computation of the spatio-temporal paths in a system. Specifically, we extract IDs and ODs from the unweighted static aggregate of the original temporal network.

3.3. Vulnerability measures

The analysis of robustness depends on the indicators used to quantify network performance. Here we introduce the measures we use to evaluate the topological, temporal and spatial vulnerability of a spatio-temporal network. While there are many properties relevant to the study of the function of a network, we are careful to select measures that do not confound these three properties.

An important property of real-world networks is the size of the largest strongly connected component, as this represents the extent to which the network is mutually reachable. This is commonly used to study topology in static and temporal contexts [9,12,44]. The notion of a connected component in a spatio-temporal network follows from the temporal definition in [53]. Formally, a strongly connected component A is a set of nodes where there exists a spatio-temporal path between all pairs of nodes $v, w \in A$. The *giant strong component size* S is the size of the largest strongly connected component as a fraction of the overall number of nodes N . The definition of *giant weak component size* follows similarly. However, we focus our robustness analysis on strong component size as this measure is especially relevant in many real-world systems; for example, the resilience of the giant strong component represents the ability of a public transport system to retain mutual navigability when subject to station closures. We note that in this paper we deal with giant components in finite-size real-world systems.

Efficiency [59,60] gives us a method to explore the overall spatial and temporal performance of a network. For a chosen distance measure (i.e. temporal λ or spatial ζ), efficiency tells us how effective a network is in supporting rapid transmission with minimum distance travelled. Formally, *temporal efficiency* (first defined in [61]) is the average reciprocal latency over all pairs of nodes in the network, expressed as

$$\mathcal{E}^\lambda = \frac{1}{N(N-1)} \sum_{v \neq w} \frac{1}{\text{dist}_{vw}^\lambda}. \quad (3.5)$$

This quantity is normalized between 0 and 1. In our spatio-temporal framework, $\mathcal{E}^\lambda = 1$ is achieved if direct propagation between all pairs of nodes occurs in one timestep. On the other hand, $\mathcal{E}^\lambda = 0$ if propagation does not succeed between any nodes, which can occur if there are no transmission links, or propagation speeds are insufficient to complete direct propagation between nodes.

Our examination of network vulnerability focuses on the system's response to the failure of one or more nodes, which has previously been measured in temporal networks using temporal robustness [40]. In this paper, we study the effect of complete deactivation of a node; i.e. where the failed node no longer has any connections to or from the rest of the system. Consider a set of deactivated nodes $D \subseteq V$. We denote the temporal efficiency of a network with deactivated nodes D by $\mathcal{E}^\lambda(D)$. The *relative efficiency* is then given by $\mathcal{E}^\lambda(D)/\mathcal{E}^\lambda$. When evaluating a systematic attack, the choice of nodes D depends on the chosen attack strategy. An effective strategy is one that can cause significant damage with few node deactivations, thus we are particularly interested in the relative efficiency with respect to the fraction f of nodes deactivated, $f = |D|/N$. Formally, we refer to this as the network's *temporal robustness* $R^\lambda(f)$ after the fraction f of nodes have been deactivated. In probabilistic failure models, such as uniform random deactivation, $R^\lambda(f)$ represents the expected relative temporal efficiency given failure probability f . (In general, measuring robustness through the relative decline in a global network quantity is sometimes referred to as the R -value of that quantity [62].)

While temporal efficiency measures the latency of the network, its spatial counterpart, *spatial efficiency* \mathcal{E}^ζ , is calculated over the reciprocal spatial distances between nodes in the network. Node failure will force paths to take alternative routes in the network, typically over longer physical distances, therefore reducing spatial efficiency. We denote a network's *spatial robustness* with respect to a given deactivation rate f by $R^\zeta(f)$.

3.4. Computation of attack strategies and measures

In terms of computation, TC, PB and BE attacks all require calculation of spatio-temporal shortest paths in the underlying network (see §2.3 for more information). Similarly, measurement of efficiency depends on the shortest distances between each pair of node in the network. Owing to the non-transitive and non-symmetric nature of spatio-temporal paths, in practice, we must use the affine graph method to compute giant component sizes (see [53] for details), which is the basis for the giant strong component size vulnerability measure. The two local-knowledge attack strategies (ID and OD) are simpler to compute, requiring only knowledge of the aggregate degree distribution.

4. Empirical spatio-temporal networks

In this paper, we analyse the spatio-temporal robustness of transport, biological and social systems through six real-world networks. Here we detail how each network is constructed. Specific choices of temporal granularity, number of snapshots and observation duration for each network can be found in table 2. The table also includes descriptive summaries of each network's properties. Further detail on the data materials used in the preparation of each network can be found in the electronic supplementary material. This also includes visualizations of the time-evolution of each network.

London Underground (London Metro). The London Underground is a rapid transit transport system that covers much of Greater London. In our spatio-temporal construction of the Underground system, we model the transits of passenger trains moving between each station along their timetabled routes. Each node is a metro station. We build a fine-grained time-varying representation of station-to-station transit speeds from the vehicles' scheduled journey times and distances. In a given snapshot, we obtain the propagation speed between two consecutive stops by averaging the speeds of trains serving those stations during the corresponding time interval. We select the timetable of February 2015 and set the observation start time to Monday at 00.00.

Paris Metro (Paris Metro). We obtain a spatio-temporal construction of the Paris Metro using the same approach as for London Metro, resulting in a network consisting of 302 stations. We select the timetable of December 2015 and set the observation start time to Monday at 00.00.

New York Subway (New York Metro). Our spatio-temporal construction of the New York City Subway also follows the same approach as London Metro. We select the weekday timetable of December 2015 and an observation start time of Monday 00.00. Multiple stations comprising the same station complex are unified into one node, yielding 417 subway stations active in the observation period.

US Domestic Flights (US Flights). The flights network is constructed from actual take-off and landing times of domestic passenger flights in the USA in the month of February 2014. Each node is a US airport. Our approach to extracting transit speeds is similar to London Metro. Time zones are normalized to EST, and we start our observation window at Monday 00.00. Airports are spatially embedded on the WGS-84 ellipsoid and physical distances are calculated using Vincenty's equation. We note that, unlike the London Metro dataset, the flight times are from reported data, rather than *a priori* timetables. Thus, the flight durations reflect environmental conditions that may delay or hasten transits.

C. Elegans (*C. Elegans*). *Caenorhabditis elegans* (*C. Elegans*) is a nematode and the first organism to have its entire cellular nervous system mapped [63], including its growth from embryogenesis [64,65], spatial configuration [66] and wiring [5,67]. In our spatio-temporal construction of the *C. Elegans* network, we study signalling among the neurons in the adult worm. Neuron spatial coordinates were collected in [66] and are given in two-dimensions along the worm's lateral plane. Missing neuron coordinates were supplemented using the approach described in [64]. The topological structure of the system is constructed using the synaptic connection data compiled in [5,67]. The connectivity data include the type (chemical or electrical) and density of directed synaptic connections, allowing us to infer transmission delay between neurons. The task of empirically measuring synaptic latency in *C. Elegans* is particularly challenging and so far only a small number of synapses have been examined. State-of-the-art optogenetic methods show that synaptic delay over chemical-only synapses is of the order of milliseconds [68], and specifically in the range 10–30 ms (see sup. fig. 1 in [68]). We encode propagation speeds in our network according to this range of baseline measurements, scaling speed values based on the synapse type (chemical, electrical or both) and physical distance. Thus, we differentiate between these two types of connection (i.e. electrical and chemical).

Dartmouth StudentLife Experiment (StudentLife). The StudentLife experiment [69] used continuous smartphone monitoring to follow a cohort of students at Dartmouth College over one academic semester. To study the information sharing potential of this network using our spatio-temporal framework, we construct the patterns of communication between students via their SMS and phone logs. Each node corresponds to a student participating in the experiment. Communication events between phones correspond to opportunities for instantaneous information transmission between students. Contact between two students in a particular timestamp maps to a link with infinite propagation speed. In particular, calls are represented as bidirectional edges and SMS messages are treated as directed from sender to recipient. Students' mobility patterns are extracted from their GPS and wireless access point geolocation logs. Our observation window covers six weeks from Monday 8 April 2013.

Our framework offers an alternative approximation of passenger transit (i.e. in London, Paris, New York and US Flights) to other, space-agnostic temporal models. They capture the speed at which passengers can be conveyed, as well as dependence on time-ordering. We note that the aggregated

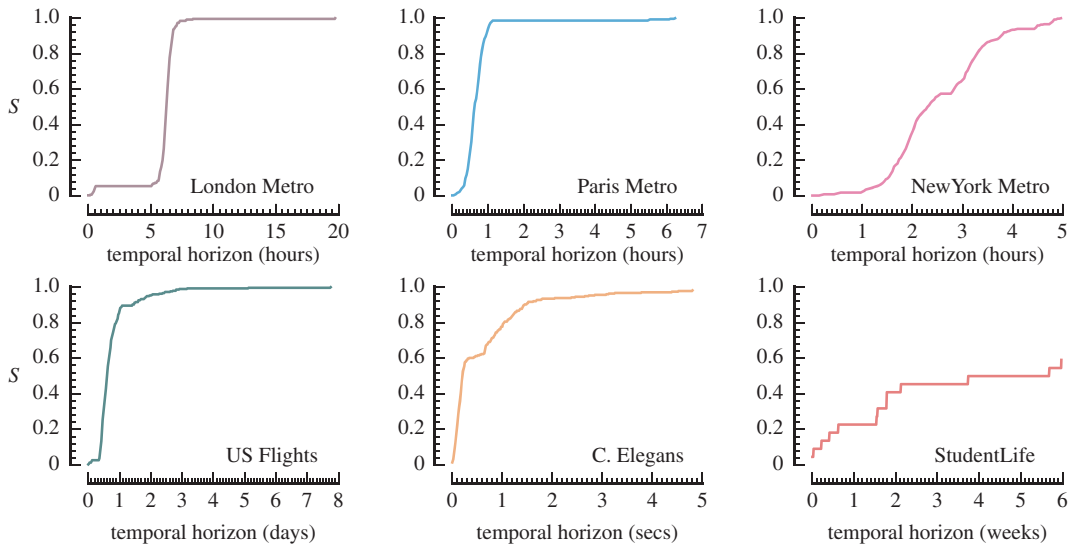


Figure 3. Influence of temporal horizon on the giant (strong) component size S . The temporal horizon h limits the observation window during which spatio-temporal paths are extracted to $(t_i, t_i + h)$.

Table 2. Summary of the six spatio-temporal systems explored in this paper: timetabled London Underground transits (London Metro), Paris Metro transits (Paris Metro), New York City Subway transits (New York Metro), US domestic flights (US Flights), the nervous system of *C. Elegans* (*C. Elegans*) and mobile communications among university students (StudentLife). For each network, we show the number of nodes N , number of directed edges $|E^*|$ in the aggregate static graph, temporal granularity τ (also referred to as window size), number of timesteps T , temporal duration, topological temporal correlation C [56], topological reciprocity \bar{r} , weight reciprocity $\bar{\rho}$, average radius of gyration \bar{y}_{gyr} over all nodes, and median propagation speed over all edges. See appendices for definitions of reciprocity (appendix A) and radius of gyration (appendix B) in a spatio-temporal system.

network	N	$ E^* $	τ	T	duration	C	\bar{r}	$\bar{\rho}$	\bar{y}_{gyr}	Med. speed
London Metro	270	628	2 min	1440	2 days	0.526	0.77	0.71	0.00 m	8.39 m s^{-1}
Paris Metro	302	705	2 min	1440	2 days	0.444	0.76	0.68	0.00 m	6.56 m s^{-1}
New York Metro	417	1058	2 min	1440	2 days	0.347	0.50	0.48	0.00 m	7.09 m s^{-1}
US Flights	299	3947	30 min	480	10 days	0.382	0.42	0.38	0.00 m	152.81 m s^{-1}
<i>C. Elegans</i>	279	2990	10 ms	1200	12 s	1.000	0.47	0.45	0.00 m	0.44 mm s^{-1}
StudentLife	22	68	60 min	1008	42 days	0.007	0.78	0.78	18.9 km	inst.

directed versions of these two networks have over 99% reciprocity, whereas their average reciprocity in the spatio-temporal representation is much lower (table 2). This highlights the loss of information caused by time aggregation. The effect is especially pronounced in US Flights, where we see that simultaneous inbound and outbound flights between two airports are rare. We also note the variation in temporal correlation over the datasets. *C. Elegans* is the only network with a static topology, giving a correlation of 1.

Figure 3 shows the growth of reachability in each spatio-temporal network as the temporal horizon is extended. In the case of the two transport networks, we see the formation of spatio-temporal paths during normal operating hours, contrasting with little or no growth during early morning. New York is unlike the other two metro systems as it operates 24/7, with a slightly reduced night service between midnight and 6.00. We also note the effect of different temporal and spatial scales present in each network. The Underground reaches full coverage within a day, whereas neural network communication is of the order of seconds. Although most of the US flight network is mutually reachable within 1 day, a number of remote off-mainland airports with limited service are responsible for delaying full coverage by up to a week. Owing to the presence of five isolated nodes in the reconstruction of the *C. Elegans* nervous system we observe that the giant strong component in *C. Elegans* does not reach full coverage. This is also observed in aggregate analyses of *C. Elegans* [67].

5. Results and discussion

In this section, we use the framework outlined in §2 to compare the response of real-world spatio-temporal networks (described in §4) when subject to the node failure models presented in §3. Before considering systematic attack strategies, we first explore the response of each network to Err.

5.1. Tolerance to random error

Figure 4 shows the response of each robustness measure with respect to uniform random failure. We used 1000 random realizations at each deactivation rate to obtain expected values for each quantity. The London network is substantially less robust according to all three measures. In figure 4*a*, we see that the network becomes highly fragmented after very few failures, and the giant component is effectively eliminated (filling less than 5% of the network) at failure rates above 0.6.

In an ideal configuration, deactivating one node should have minimum effect in disconnecting other nodes in the network. We see that US Flights and C. Elegans are much closer to this resilient behaviour. When comparing the edge density of the directed aggregate networks (0.9% for London, 0.7% for Paris, 0.6% for New York, 4.4% for US Flights, 3.9% for C. Elegans and 14.7% for StudentLife) this may in part be explained by more path redundancy in the high-density networks. Although the giant component size in the intact StudentLife network is smaller, the profile of its relative decay is similar to that of US Flights and C. Elegans. We also observe that C. Elegans is typically more resilient than US Flights, despite C. Elegans having slightly lower density. C. Elegans also exhibits increasing giant-component vulnerability in the $f = 0.5\text{--}0.8$ range, eventually becoming more fragmented than US Flights. Its temporal and spatial robustness (figure 4*b,c*), on the other hand, do not exhibit the same behaviour.

Comparing figure 4*b,c*, we note that the spatial robustness of US Flights, C. Elegans, London and Paris follow a degradation pattern similar to their temporal robustness. In these systems, spatio-temporal paths are constrained by finite node-to-node propagation speeds. The propagation speeds in these networks are also heterogeneous, with the amount of diversity depending on the particular system; for example, longer track segments in London and longer flight paths in US Flights tend to have higher average speeds (flight duration also depends on atmospheric conditions), and synaptic transmission speeds in C. Elegans depend on the signalling mechanism. On the other hand, the homogeneous infinite propagation speeds in StudentLife permit transmission to occur instantaneously, independent of the physical separation between individuals, and thus we see very different behaviour in this network with respect to spatial and temporal robustness (figure 4*b,c*).

5.2. Static spatial paths versus spatio-temporal paths

Of the three urban transit systems, we note that New York exhibits the greatest difference in response to random failure between spatial and temporal robustness (figure 4*b* versus *c*; area under curve comparison in table 3). In figure 5, we demonstrate the differing routes taken according to spatio-temporal paths versus shortest paths in the equivalent aggregate spatial network. We find that spatio-temporal distances tend to take longer routes through the network than purely spatial paths in New York. Indeed, correlation between these pairwise distances is least in the New York network; that is, minimum-latency routes tend to traverse longer distances in the New York transit system. We note that the comparison in this figure is in the intact networks (i.e. without node deactivations), to highlight the differences in the types of paths in these systems.

5.3. Tolerance to systematic attack

To study attack tolerance, we apply the deactivation strategies defined in §3 to each network, allowing us to measure network performance after a fraction of nodes f are deactivated.

Although all attack schemes are generally more effective than Err, there is substantial variation in their performance. In figure 6, we plot the effectiveness of each strategy at dismantling the giant component. In US Flights, the damage caused by the four non-closeness attacks is very similar. We see some slight variation in giant-component vulnerability with respect to each attack (figure 6*d*), and very similar behaviour in temporal vulnerability (figure 7*d*). This contrasts with London and C. Elegans, where the betweenness-based attacks follow substantially different profiles to the degree-based attacks.

The two betweenness attacks are very effective at dismantling the giant component in London up to $f = 0.05$ (figure 6*a*). At this point (i.e. with 5% of the nodes deactivated), the nodes vital to connecting

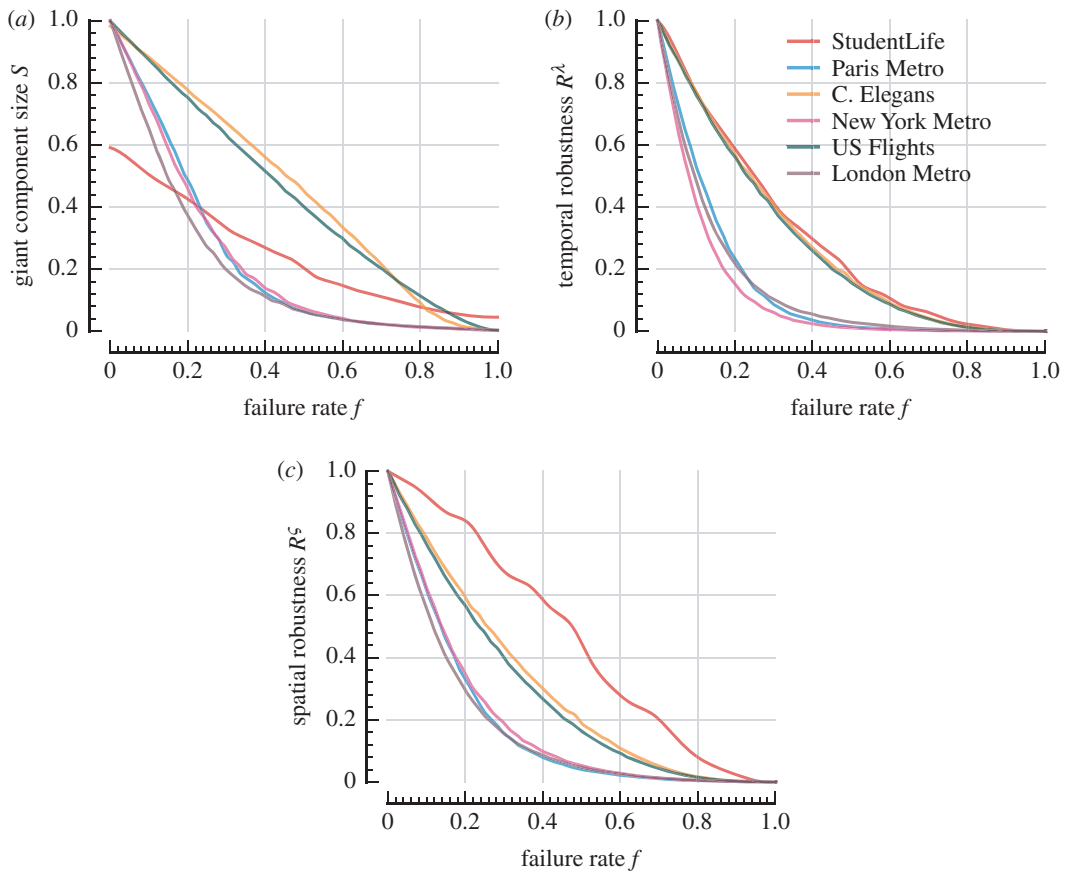


Figure 4. Vulnerability to random node failure measured by giant strongly connected component size S , temporal robustness R^λ , spatial robustness R^σ . Expected values for quantities were obtained through simulation of 1000 random realizations at each deactivation rate and are represented by solid curves. Standard error (shaded region surrounding each curve) is negligible. (a) Giant component size, (b) temporal robustness and (c) spatial robustness.

Table 3. Overall effectiveness of PB attack and BE attack measured by the area under each robustness curve. Areas are provided for two vulnerability measures: giant strong temporal component size (S) and temporal robustness (R^λ). A lower value indicates an attack that is more aggressive. Random failure (Err) is also included for comparison. Values in bold indicate the most-aggressive of the two attack strategies shown.

	area under robustness curve					
	S			R^λ		
	Err	PB	BE	Err	PB	BE
London Metro	0.195	0.035	0.039	0.134	0.042	0.034
Paris Metro	0.224	0.043	0.043	0.136	0.030	0.030
New York Metro	0.221	0.031	0.031	0.110	0.024	0.022
US Flights	0.431	0.060	0.060	0.274	0.025	0.025
C. Elegans	0.448	0.201	0.234	0.278	0.124	0.106
StudentLife	0.244	0.132	0.152	0.287	0.202	0.190

London's central core to peripheral Underground lines have been deactivated. A substantial proportion of nodes (roughly 190 of 270 stations) belong to isolated clusters in this now-disconnected outer region. There is more tolerance to failure in the remaining core, as we observe from the gradual decline in S from $f = 0.05$, and the giant component is finally eliminated at $f = 0.45$. Interestingly, the threshold at which




Figure 5. Three urban transport networks and the US flights network. Comparison of spatio-temporal paths to shortest spatial paths in the aggregate spatial network, measured by spatial distance. Shortest spatial paths follow minimum spatial length. Each heatmap compares the spatial distances between pairs of stations according to definitions of spatial path and spatio-temporal path. Pearson correlation between these distances is 0.983, 0.952, 0.871 and 0.950 for London, Paris, New York and US Flights, respectively. (a) London Metro, (b) Paris Metro, (c) New York Metro and (d) US Flights.

systematic attack eliminates the giant component in US Flights (figure 6d) is much lower; specifically, this occurs at $f = 0.16$ for US Flights and at $f = 0.45$ for London.

Comparing figures 6 and 7, we observe how the two betweenness centralities attack reachability and temporal efficiency differently. PB preferentially attacks nodes through which many geodesic paths flow, whereas BE prioritizes nodes supporting temporally efficient paths. Thus, the giant component is harmed more rapidly in the PB attack (figure 6), and temporal robustness declines more quickly under the BE attack (figure 7). This difference is especially pronounced in *C. Elegans* (figures 6e and 7e).

To judge the overall effectiveness of a particular attack for a given performance measure, we also consider the area under each robustness curve. Values for random failure and the two betweenness-based attacks are provided in table 3. We see that these two attacks indeed preferentially degrade different features of a network. Specifically, table 3 highlights that for London, *C. Elegans* and StudentLife, the lowest overall S -robustness is achieved with the PB attack strategy and the lowest overall R^λ -robustness is achieved with the BE attack strategy. We also find that in the other three networks (US Flights, New York and Paris), the effect is much less pronounced, where both PB and BE have roughly the same effect with respect to the two vulnerability measures (S and R^λ).

In general, we observe that both betweenness-based attacks harm each network more effectively than the degree-based strategies. This is intuitive, as betweenness-based attacks are able to leverage global network information to decide their targets. What is more surprising is that TC, which is also a




Figure 6. Robustness of the giant strongly connected component when each system is subject to five systematic attack strategies: TC, PB, BE, ID and OD. Vulnerability to random failure (Err) included for comparison. Shaded regions represent standard deviation (see figure 4 for standard error). (a) London Metro, (b) Paris Metro, (c) New York Metro, (d) US Flights, (e) C. Elegans and (f) StudentLife.




Figure 7. Robustness of temporal efficiency under systematic attack. Labels and colours same as figure 6. Shaded regions represent standard deviation (see figure 4 for standard error). (a) London Metro, (b) Paris Metro, (c) New York Metro, (d) US Flights, (e) C. Elegans and (f) StudentLife.

global-knowledge attack, is ineffective at identifying vulnerabilities in the network, and is often unable to out-perform the naive degree-based attacks. The same is true for the strategy's effectiveness in attacking the network's temporal efficiency (figure 7). Temporal (in-)closeness specifically identifies 'sink' nodes that are effective endpoints for information flow. These results show that such sinks do not tend to be

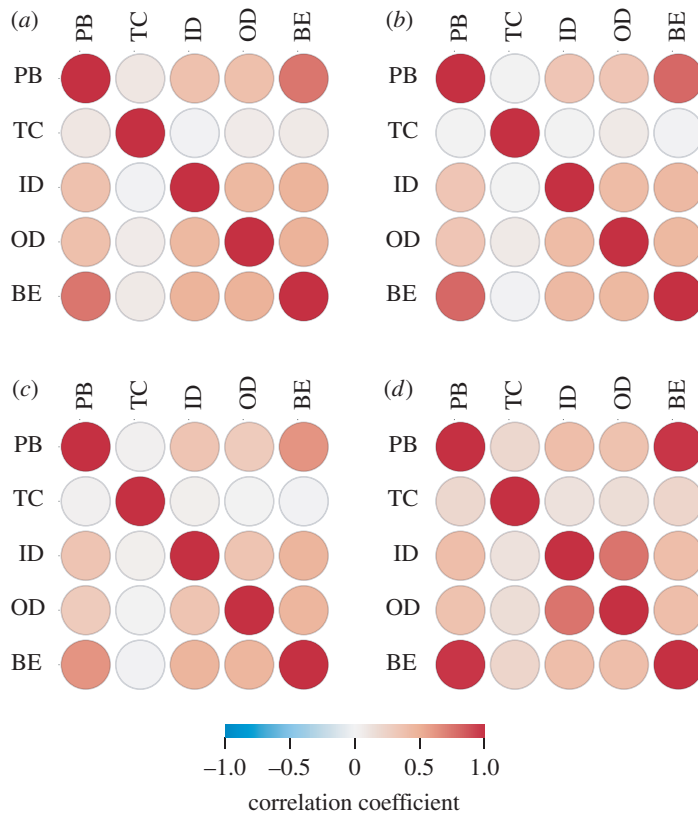


Figure 8. Three urban transport networks and the US flights network. Comparison of the node deactivation order according to each strategy in each network. Correlation is measured by the Kendall rank correlation coefficient (Kendall's τ). Attack strategies as follows: PB, TC, ID, OD and BE. (a) London Metro, (b) Paris Metro, (c) New York Metro and (d) US Flights.

important centres on which other nodes rely, and therefore their deactivation has relatively little impact on the performance of the network.

We see that *C. Elegans* is substantially more robust to systematic attack. Using the most-effective strategy in each network as a comparison, we see the giant component in *C. Elegans* is highly resilient (figure 6e), managing to remain at over 80% coverage up to $f = 0.12$, whereas other networks begin rapid degradation at $f = 0.01$ or lower.

Finally, we compare the responses of the four transport networks to the most-aggressive attack strategies. We see that temporal and giant component degradation for London (figures 6a and 7a, respectively) both follow very similar patterns, indicating that mutual reachability declines at a similar rate to the relative decline in temporal efficiency. This contrasts with US Flights, whose giant component (figure 6d) is more resilient than its temporal robustness (figure 7d). This differing behaviour shows that there is a period during systematic attack on US Flights where the network suffers increasing temporal delay while its giant component remains almost intact.

5.4. Attack strategy correlations

To explore the attack strategies further, we compare the order with which each strategy preferentially deactivates nodes (figure 8). Unsurprisingly, examples where we see the biggest difference in attack performance are also where we see lowest correlation in deactivation order. For example, in Paris Metro and US Flights we see high correlation among most pairs of strategies, and also observe that these strategies have similar robustness profiles, especially in the case of US Flights. Comparing centrality correlation between the two betweenness attacks (BE and PB) also offers insight. These are overall highly correlated in all networks. Specifically, we observe correlation coefficients 0.732 for London Metro, 0.788 for Paris Metro, 0.613 for New York Metro, 0.961 for US Flights and 0.613 for *C. Elegans* (not plotted). However, where we see less agreement between these two strategies, we also see the greatest differences in their effectiveness at attacking the giant component (figure 6) versus temporal robustness (figure 7).




Figure 9. Comparison of spatio-temporal attacks (PB and BE) and the pure-temporal path betweenness attack (PTPB). (a) London Metro and (b) New York Metro.

This effect is most pronounced in *C. Elegans* and New York Metro, which also have the lowest PB–BE correlation of the networks we studied.

5.5. Comparison with temporal centrality

Finally, we compare the two attack strategies based on spatio-temporal paths (PB and BE) to a purely temporal counterpart. Here we explore whether a centrality measure based on temporally shortest paths has the same effect as PB and BE. Specifically, we apply the temporal betweenness centrality introduced in [50], which for convenience we refer to as PTPB (pure temporal path betweenness). In the calculation of temporal betweenness centrality, a node’s betweenness is counted as the number of shortest temporal paths through it, as opposed to the number of shortest spatio-temporal paths. In other words, these temporal paths are not constrained by distance and propagation speed. (Equivalently, we can regard this as the case where all propagation speeds are infinite.) We show the performance of these attacks for two metro systems, London Metro and New York Metro, in figure 9. We see that the space-agnostic attack is less effective than PB and BE in both systems.

6. Conclusion

The ability to explicitly represent non-instantaneous propagation that reflects the space and time in which a network is embedded allows us to better understand the properties of real-world spatio-temporal systems. In this paper, we develop a framework that combines these two dimensions, constructing networks whose connectivity is constrained by the dynamics of the real-world systems they represent. With this framework, we explore the resilience of these systems in terms of their topological, spatial and temporal behaviour. The proposed approach is based on a constrained-propagation model over a time-varying representation of interaction speeds between nodes and is a re-formulation of previous discrete-time instantaneous-spreading models of temporal networks. Through numerical experiments on empirical systems, we study the spatial, temporal and topological performance of spatio-temporal networks when subject to Err and systematic attack. Our approach is particularly motivated by the joint spatial and temporal nature of urban transport systems, while we additionally also test the generality of the approach on other classes of network. We introduce global- and local-knowledge attack

strategies, finding that attacks based on the betweenness of nodes perform better than other strategies. Most saliently, we find that the spatial, temporal and topological responses of a system to node failure can behave differently to the same attack strategy. In particular, we find that systematic attack based on betweenness centrality is typically most effective at harming the temporal efficiency of a network, leading to increased delays, whereas PB is most effective at eliminating mutual reachability within a network.

Data accessibility. Datasets [70] used in this research are available via the Dryad Repository (<http://dx.doi.org/10.5061/dryad.3p27r>). Visualizations [71] of these networks are available via the Figshare Repository (<https://dx.doi.org/10.6084/m9.figshare.1452976>). These datasets and visualizations are further described in this paper's electronic supplementary material.

Authors' contributions. M.J.W. and M.M. designed and developed the research, performed the analyses, reviewed the results and wrote the manuscript.

Competing interests. The authors declare that they have no competing interests.

Funding. This work was supported by LASAGNE (Multi-layer Spatiotemporal Generalized Networks), a European Commission STREP project (grant no. 318132), and *The Uncertainty of Identity*, an EPSRC project (grant no. EP/J005266/1).

Acknowledgements. The authors would like to thank Fani Tsapeli and Luca Canzian for their help in preparing the StudentLife mobility data, Petra Vertes for insights in reconstructing the *C. Elegans* connectome, Manuel Zimmer for advice on synaptic delay and Luca Rossi for fruitful discussions. Data processing and numerical experiments were performed in part using the computational facilities of the University of Birmingham's BEAR division.

Appendix A. Average reciprocity in a time-varying network

The reciprocity of a (single) directed graph characterizes the overall extent to which each edge in the network is in some way balanced by an edge in the opposite direction. This measure is commonly defined for binary graphs [72] and has more recently been adapted to study weighted graphs [73]. Here we describe how we extend this measure to characterize the average reciprocity in a time-varying weighted directed network represented by the sequence of weight matrices

$$\mathbf{s}^{[t_1]}, \mathbf{s}^{[t_2]}, \dots, \mathbf{s}^{[t_T]}. \quad (\text{A } 1)$$

First, let us consider a single weight matrix $\mathbf{s}^{[t]}$ and obtain the global weighted reciprocity $\rho^{[t]}$ of this (static) graph as defined by Squartini *et al.* [73]. That is, given the total weight of the graph

$$W^{[t]} = \sum_v \sum_{w \neq v} \mathbf{s}_{vw}^{[t]} \quad (\text{A } 2)$$

and the total reciprocated weight of the graph

$$W_{\leftrightarrow}^{[t]} = \sum_v \sum_{w \neq v} \min(\mathbf{s}_{vw}^{[t]}, \mathbf{s}_{wv}^{[t]}), \quad (\text{A } 3)$$

we obtain the global weighted reciprocity as

$$\rho^{[t]} = \frac{W_{\leftrightarrow}^{[t]}}{W^{[t]}}. \quad (\text{A } 4)$$

In the case that the graph is empty, we set $\rho^{[t]} = 1$. Finally, as a diagnostic of the reciprocity in the whole time-varying network, we define the *weight reciprocity* as the average weighted reciprocity

$$\bar{\rho} = \frac{1}{T} \sum_{i=1}^T \rho^{[t_i]}. \quad (\text{A } 5)$$

Weight reciprocity is normalized between 0 and 1. $\bar{\rho} = 1$ indicates that each edge in the network is reciprocated by an identically weighted edge in the opposite direction, and thus the network is effectively undirected. In the opposite case, where each edge in the network has no reciprocal counterpart, we have $\bar{\rho} = 0$. We note that for the trivial case of an empty (edgeless) time-varying network we have $\bar{\rho} = 1$.

It is sometimes useful to ignore weights in the time-varying network and study reciprocity from a purely topological viewpoint. We refer to this as the (average) *topological reciprocity*, denoted by \bar{r} . This quantity is the average reciprocity over the network's sequence of binary adjacency matrices.

Appendix B. Radius of gyration

To characterize the typical physical distance travelled by a node in a spatio-temporal system, we use the average radius of gyration taken over all nodes in the system. We follow the same approach as previous work that has adapted the classical mechanical definition of radius of gyration to discrete-time trajectories of moving objects [74].

Recall that in our representation of a spatio-temporal system, nodes are embedded in a k -dimensional metric space with physical distance function g . Given a node v and its trajectory represented by the time-varying sequence of positions $l_v^{[1]}, l_v^{[2]}, \dots, l_v^{[T]}$ the radius of gyration of v is expressed as

$$y_{\text{gyr}}^v = \left(\frac{1}{T} \sum_{i=1}^T g(l_v^{[i]}, y_{\mu}^v)^2 \right)^{1/2}, \quad (\text{B } 1)$$

where y_{μ}^v is the centroid of the trajectory of v . Our summary of the mobility of the set of all nodes V in a system is then given by the average radius of gyration

$$\bar{y}_{\text{gyr}} = \frac{1}{|V|} \sum_{v \in V} y_{\text{gyr}}^v. \quad (\text{B } 2)$$

References

- Newman M. 2010 *Networks: an introduction*. Oxford, UK: Oxford University Press.
- Newman MEJ, Barabási A-L, Watts DJ. 2006 *The structure and dynamics of networks*. Princeton, NJ: Princeton University Press.
- Gallotti R, Barthélemy M. 2015 The multilayer temporal network of public transport in Great Britain. *Sci. Data* **2**, 140056. (doi:10.1038/sdata.2014.56)
- Bota M, Sporns O, Swanson LW. 2015 Architecture of the cerebral cortical association connectome underlying cognition. *Proc. Natl Acad. Sci. USA* **112**, E2093–E2101. (doi:10.1073/pnas.1504394112)
- Chen BL, Hall DH, Chklovskii DB. 2006 Wiring optimization can relate neuronal structure and function. *Proc. Natl Acad. Sci. USA* **103**, 4723–4728. (doi:10.1073/pnas.0506806103)
- Yook S-H, Jeong H, Barabási A-L. 2002 Modeling the Internet's large-scale topology. *Proc. Natl Acad. Sci. USA* **99**, 13382–13386. (doi:10.1073/pnas.172501399)
- Solé RV, Rosas-Casals M, Corominas-Murtra B, Valverde S. 2008 Robustness of the European power grids under intentional attack. *Phys. Rev. E* **77**, 026102. (doi:10.1103/PhysRevE.77.026102)
- Eagle N, Pentland A, Lazer D. 2009 Inferring friendship network structure by using mobile phone data. *Proc. Natl Acad. Sci. USA* **106**, 15 274–15 278. (doi:10.1073/pnas.0900282106)
- Albert R, Jeong H, Barabási A-L. 2000 Error and attack tolerance of complex networks. *Nature* **406**, 378–382. (doi:10.1038/35019019)
- Cohen R, Erez K, ben Avraham D, Havlin S. 2000 Resilience of the internet to random breakdowns. *Phys. Rev. Lett.* **85**, 4626–4628. (doi:10.1103/PhysRevLett.85.4626)
- de Domenico M, Solé-Ribalta A, Gómez S, Arenas A. 2014 Navigability of interconnected networks under random failures. *Proc. Natl Acad. Sci. USA* **111**, 8351–8356. (doi:10.1073/pnas.1318469111)
- Holme P, Kim BJ, Yoon CN, Han SK. 2002 Attack vulnerability of complex networks. *Phys. Rev. E* **65**, 056109. (doi:10.1103/PhysRevE.65.056109)
- Callaway DS, Newman MEJ, Strogatz SH, Watts DJ. 2000 Network robustness and fragility: percolation on random graphs. *Phys. Rev. Lett.* **85**, 5468–5471. (doi:10.1103/PhysRevLett.85.5468)
- Albert R, Albert I, Nakarado GL. 2004 Structural vulnerability of the North American power grid. *Phys. Rev. E* **69**, 025103. (doi:10.1103/PhysRevE.69.025103)
- Berezin Y, Bashan A, Danziger MM, Li D, Havlin S. 2015 Localized attacks on spatially embedded networks with dependencies. *Sci. Rep.* **5**, 08934. (doi:10.1038/srep08934)
- Buldirev SV, Parshani R, Paul G, Stanley HE, Havlin S. 2010 Catastrophic cascade of failures in interdependent networks. *Nature* **464**, 1025–1028. (doi:10.1038/nature08932)
- Berche B, von Ferber C, Holovatch T, Holovatch Y. 2009 Resilience of public transport networks against attacks. *Eur. Phys. J. B* **71**, 125–137. (doi:10.1140/epjb/e2009-00291-3)
- Dall'Asta L, Barrat A, Barthélemy M, Vespignani A. 2006 Vulnerability of weighted networks. *J. Stat. Mech. Theor. Exp.* **2006**, P04006. (doi:10.1088/1742-5468/2006/04/P04006)
- Fornito A, Zalesky A, Breakspear M. 2015 The connectomics of brain disorders. *Nat. Rev. Neurosci.* **16**, 159–172. (doi:10.1038/nrn3901)
- Stam CJ. 2014 Modern network science of neurological disorders. *Nat. Rev. Neurosci.* **15**, 683–695. (doi:10.1038/nrn3801)
- Helbing D. 2013 Globally networked risks and how to respond. *Nature* **497**, 51–59. (doi:10.1038/nature12047)
- Thurner S, Poledna S. 2013 DebtRank-transparency: controlling systemic risk in financial networks. *Sci. Rep.* **3**, 01888. (doi:10.1038/srep01888)
- Schneider CM, Mihaljev T, Havlin S, Herrmann HJ. 2011 Suppressing epidemics with a limited amount of immunization units. *Phys. Rev. E* **84**, 061911. (doi:10.1103/PhysRevE.84.061911)
- Holme P, Saramäki J. 2012 Temporal networks. *Phys. Rep.* **519**, 97–125. (doi:10.1016/j.physrep.2012.03.001)
- Gallotti R, Barthélemy M. 2014 Anatomy and efficiency of urban multimodal mobility. *Sci. Rep.* **4**. (doi:10.1038/srep06911)
- Angeloudis P, Fisk D. 2006 Large subway systems as complex networks. *Physica A* **367**, 553–558. (doi:10.1016/j.physa.2005.11.007)
- Lee K, Jung W-S, Park JS, Choi MY. 2008 Statistical analysis of the metropolitan Seoul subway system: network structure and passenger flows. *Physica A* **387**, 6231–6234. (doi:10.1016/j.physa.2008.06.035)
- Derrible S, Kennedy C. 2010 The complexity and robustness of metro networks. *Physica A* **389**, 3678–3691. (doi:10.1016/j.physa.2010.04.008)
- Roth C, Kang SM, Batty M, Barthélemy M. 2012 A long-time limit for world subway networks. *J. R. Soc. Interface* **9**, 2540–2550. (doi:10.1098/rsif.2012.0259)
- Barthélemy M. 2011 Spatial networks. *Phys. Rep.* **499**, 1–101. (doi:10.1016/j.physrep.2010.11.002)
- Kivela M, Arenas A, Barthélemy M, Gleeson JP, Moreno Y, Porter MA. 2014 Multilayer networks. *J. Complex Networks* **2**, 203–271. (doi:10.1093/comnet/cnu016)
- Pan RK, Saramäki J. 2011 Path lengths, correlations, and centrality in temporal networks. *Phys. Rev. E* **84**, 016105. (doi:10.1103/PhysRevE.84.016105)
- Dorbritz R, Weidmann U. 2009 Stability of public transportation systems in case of random failures and intended attacks—a case study from Switzerland. In *Proc. 4th IET Int. Conf. on System Safety, Incorporating the SaRS Annual Conf.*, pp. 1–6. London, UK: IET.
- Trajanovski S, Kuipers FA, Ilić A, Crowcroft J, Miegheem PV. 2015 Finding critical regions and region-disjoint paths in a network. *IEEE/ACM Trans. Netw.* **23**, 908–921. (doi:10.1109/TNET.2014.2309253)
- Neumayer S, Efrat A, Modiano E. 2015 Geographic max-flow and min-cut under a circular disk failure model. *Comput. Netw.* **77**, 117–127. (doi:10.1016/j.comnet.2014.10.026)
- Neumayer S, Zussman G, Cohen R, Modiano E. 2011 Assessing the vulnerability of the fiber infrastructure to disasters. *IEEE/ACM Trans. Netw.* **19**, 1610–1623. (doi:10.1109/TNET.2011.2128879)
- Agarwal PK, Efrat A, Ganjugunte SK, Hay D, Sankararaman S, Zussman G. 2013 The resilience of

- WDM networks to probabilistic geographical failures. *IEEE/ACM Trans. Netw.* **21**, 1525–1538. (doi:10.1109/TNET.2012.2232111)
38. Boccaletti S, Bianconi G, Criado R, del Genio CI, Gómez-Gardeñes J, Romance M, Sendiña-Nadal I, Wang Z, Zanin M. 2014 The structure and dynamics of multilayer networks. *Phys. Rep.* **544**, 1–122. (doi:10.1016/j.physrep.2014.07.001)
39. Korschake M, Lentz HHK, Conraths FJ, Hövel P, Selhorst T. 2013 On the robustness of in- and out-components in a temporal network. *PLoS ONE* **8**, E55223. (doi:10.1371/journal.pone.0055223)
40. Trajanovski S, Scellato S, Leontiadis I. 2012 Error and attack vulnerability of temporal networks. *Phys. Rev. E* **85**, 066105. (doi:10.1103/PhysRevE.85.066105)
41. Karsai M, Kivela M, Pan RK, Kaski K, Kertész J, Barabási A-L, Saramäki J. 2011 Small but slow world: how network topology and burstiness slow down spreading. *Phys. Rev. E* **83**, 025102. (doi:10.1103/PhysRevE.83.025102)
42. Kivela M, Pan RK, Kaski K, Kertész J, Saramäki J, Karsai M. 2012 Multiscale analysis of spreading in a large communication network. *J. Stat. Mech.* **2012**, P03005. (doi:10.1088/1742-5468/2012/03/P03005)
43. Scellato S, Leontiadis I, Mascolo C, Basu P, Zafer M. 2013 Evaluating temporal robustness of mobile networks. *IEEE Trans. Mob. Comput.* **12**, 105–117. (doi:10.1109/TMC.2011.248)
44. Sur S, Ganguly N, Mukherjee A. 2015 Attack tolerance of correlated time-varying social networks with well-defined communities. *Physica A* **420**, 98–107. (doi:10.1016/j.physa.2014.08.074)
45. Kempe D, Kleinberg J, Kumar A. 2000 Connectivity and inference problems for temporal networks. In *Proc. 32nd Symp. on Theory of Computing, STOC '00, Portland, OR*, pp. 504–513. New York, NY: ACM.
46. Holme P. 2005 Network reachability of real-world contact sequences. *Phys. Rev. E* **71**, 046119. (doi:10.1103/PhysRevE.71.046119)
47. Grindrod P, Higham DJ. 2010 Evolving graphs: dynamical models, inverse problems and propagation. *Proc. R. Soc. A* **466**, 753–770. (doi:10.1098/rspa.2009.0456)
48. Kaluza P, Kölzsch A, Gastner MT, Blasius B. 2010 The complex network of global cargo ship movements. *J. R. Soc. Interface* **7**, 1093–1103. (doi:10.1098/rsif.2009.0495)
49. Bajardi P, Barrat A, Natale F, Savini L, Colizza V. 2011 Dynamical patterns of cattle trade movements. *PLoS ONE* **6**, E19869. (doi:10.1371/journal.pone.0019869)
50. Nicosia V, Tang J, Mascolo C, Musolesi M, Russo G, Latora V. 2013 Graph metrics for temporal networks. In *Temporal networks* (eds P Holme, J Saramäki), pp. 15–40. Berlin, Germany: Springer.
51. Holme P. 2014 Analyzing temporal networks in social media. *Proc. IEEE* **102**, 1922–1933. (doi:10.1109/JPROC.2014.2361326)
52. Holme P, Saramäki J. 2013 Temporal networks as a modeling framework. In *Temporal networks* (eds P Holme, J Saramäki), pp. 1–14. Berlin, Germany: Springer.
53. Nicosia V, Tang J, Musolesi M, Russo G, Mascolo C, Latora V. 2012 Components in time-varying graphs. *Chaos* **22**, 023101. (doi:10.1063/1.3697996)
54. Fenu C, Higham DJ. 2015 Block matrix formulations for evolving networks. (<http://arxiv.org/abs/1511.07305>)
55. Alsayed A, Higham DJ. 2015 Betweenness in time dependent networks. *Chaos, Solitons Fractals* **72**, 35–48. (doi:10.1016/j.chaos.2014.12.009)
56. Tang J, Scellato S, Musolesi M, Mascolo C, Latora V. 2010 Small-world behavior in time-varying graphs. *Phys. Rev. E* **81**, 055101. (doi:10.1103/PhysRevE.81.055101)
57. Brandes U. 2001 A faster algorithm for betweenness centrality. *J. Math. Sociol.* **25**, 163–177. (doi:10.1080/0022250X.2001.9990249)
58. Tang J, Kim H, Mascolo C, Musolesi M. 2012 STOP: Socio-Temporal Opportunistic Patching of short range mobile malware. In *Proc. IEEE Int. Symp. on a World of Wireless, Mobile and Multimedia Networks*, pp. 1–9. San Francisco, CA: IEEE.
59. Latora V, Marchiori M. 2001 Efficient behavior of small-world networks. *Phys. Rev. Lett.* **87**, 198701. (doi:10.1103/PhysRevLett.87.198701)
60. Latora V, Marchiori M. 2003 Economic small-world behavior in weighted networks. *Eur. Phys. J. B* **32**, 249–263. (doi:10.1140/epjb/e2003-00095-5)
61. Tang J, Musolesi M, Mascolo C, Latora V. 2009 Temporal distance metrics for social network analysis. In *Proc. 2nd Workshop on Online Social Networks, Barcelona, Spain*, pp. 31–36. New York, NY: ACM.
62. Trajanovski S, Martín-Hernández J, Winterbach W, Miegheem PV. 2013 Robustness envelopes of networks. *J. Complex Networks* **1**, 44–62. (doi:10.1093/comnet/cnt004)
63. White JG, Southgate E, Thomson JN, Brenner S. 1986 The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Phil. Trans. R. Soc. Lond. B* **314**, 1–340. (doi:10.1098/rstb.1986.0056)
64. Varier S, Kaiser M. 2011 Neural development features: spatio-temporal development of the *Caenorhabditis elegans* neuronal network. *PLoS Comput. Biol.* **7**, E1001044. (doi:10.1371/journal.pcbi.1001044)
65. Nicosia V, Vértes PE, Schafer WR, Latora V, Bullmore ET. 2013 Phase transition in the economically modeled growth of a cellular nervous system. *Proc. Natl Acad. Sci. USA* **110**, 7880–7885. (doi:10.1073/pnas.1300753110)
66. Kaiser M, Hilgetag CC. 2006 Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. *PLoS Comput. Biol.* **2**, E95. (doi:10.1371/journal.pcbi.0020095)
67. Varshney LR, Chen BL, Paniagua E, Hall DH, Chklovskii DB. 2011 Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS Comput. Biol.* **7**, E1001066. (doi:10.1371/journal.pcbi.1001066)
68. Lindsay TH, Thiele TR, Lockery SR. 2011 Optogenetic analysis of synaptic transmission in the central nervous system of the nematode *Caenorhabditis elegans*. *Nat. Commun.* **2**, 306. (doi:10.1038/ncomms1304)
69. Wang R, Chen F, Chen Z, Li T, Harari G, Tignor S, Zhou X, Ben-Zeev D, Campbell AT. 2014 StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proc. 2014 Conf. on Ubiquitous Computing, UbiComp '14, Seattle, WA*, pp. 3–14. New York, NY: ACM.
70. Williams MJ, Musolesi M. 2016 Data from: spatio-temporal networks: reachability, centrality, and robustness. *Dryad Digital Repository*. (doi:10.5061/dryad.3p27r)
71. Williams MJ, Musolesi M. 2016 Visualisations from: spatio-temporal networks: reachability, centrality, and robustness. *Figshare Digital Repository*. (doi:10.6084/m9.figshare.1452976)
72. Garlaschelli D, Loffredo MI. 2004 Patterns of link reciprocity in directed networks. *Phys. Rev. Lett.* **93**, 268701. (doi:10.1103/PhysRevLett.93.268701)
73. Squartini T, Picciolo F, Ruzzenenti F, Garlaschelli D. 2013 Reciprocity of weighted networks. *Sci. Rep.* **3**, 02729. (doi:10.1038/srep02729)
74. Gonzalez MC, Hidalgo CA, Barabási A-L. 2008 Understanding individual human mobility patterns. *Nature* **453**, 779–782. (doi:10.1038/nature06958)