

Processes and Events in Dynamic Geo-Networks

Antony Galton¹ and Michael Worboys²

¹ School of Engineering, Computer Science and Mathematics,
University of Exeter, Exeter EX4 4QF, UK

A.P.Galton@exeter.ac.uk

² National Center for Geographic Information and Analysis,
University of Maine, Orono, ME 04469, USA

worboys@spatial.maine.edu

Abstract. Traditional spatial information systems hold only a single state of the ‘real world’. However, geographic phenomena have not only static but dynamic characteristics. The work described in this paper contributes to the general research effort toward a generic ontology of dynamic geographic-scale phenomena and its application to the provision of modeling, analysis, and retrieval of data in a spatio-temporal GIS. These issues are addressed in this paper with reference to *dynamic geo-networks*, that is, networks embedded in a (2-dimensional) geographic space. After an introductory and motivational section, the basic ontological categories of events and states are discussed. The paper develops these ideas in the context of flows in dynamic geo-networks, and goes on to discuss the possible kinds of causal relations. The paper concludes with an overview of the results and pointers to further research directions.

1 Introduction

Traditional spatial information systems hold only a single state of the ‘real world’. This state is almost always the most recent in time for which the data were captured. Interactions with the system therefore are ‘timeless’, in that only information contained in the single state can be retrieved. However, geographic phenomena have not only static but also dynamic characteristics. Geographic information systems are now beginning to have some temporal functionality, and a *spatio-temporal information system* manages information that is both geospatially and temporally referenced. While truly spatio-temporal information systems are still in the research arena, GISs are beginning to be extended so that they can offer some practical temporal functionality.

There are many potential application domains for spatio-temporal systems, including environmental change monitoring, transportation, socioeconomic and demographic applications, health and epidemiology, multimedia, governance and administration, crisis management, and defense. In addition to these more traditional spatio-temporal application areas, the increased use of real-time, mobile and *in situ* sensors is leading to many new potential applications for spatio-temporal data models and systems.

Spatiotemporal database research can be divided into two broad categories: that dealing with change (e.g. administrative boundary evolution, environmental change) and that dealing with moving objects (vehicles, ships, people). Although movement is a kind of change, focus on movement of objects leads to a rather different emphasis in terms of the research issues. For example, with movement, themes such as construction of suitable data structures for trajectories and prediction of future positions become dominant. A recent European project focused on movement is Chorochronos [14]. The emphasis of the present work is on more general change.

The work described in this paper contributes to the general research effort toward a generic ontology of dynamic geographic-scale phenomena and its application to the provision of modeling, analysis, and retrieval of data in a spatio-temporal GIS. There is a parallel here with the early days of static GIS, and the process of construction of basic primitives to form a generic framework around which systems could be developed. While conceptual modeling methodologies such as entity-relationship analysis [5], extensions to the relational database model [6], semantic data modeling [4,17], and the object-oriented approach [2,19,3] formed the basis for all information systems, GIS needed geometric and topological types and operators that were sufficiently universal to be deployed by a diverse range of geospatial systems, but specialized enough to be useful for the GIS community. These were developed first by the GIS research community as a collection of generic spatial data types (e.g., [15,8]), or collections of classes and operations in relational [21] and object-oriented settings [7,23,22].

2 Conceptual Models for Spatio-temporal Information Systems and Motivation for the Research

In this section we briefly review examples of conceptual modeling approaches for spatio-temporal information systems, indicate their limitations, and use these to motivate our own work. While there has been considerable progress in the underlying data structures and indexes for temporal, and to a lesser extent, spatio-temporal databases (for a review, see [1]), there has been much slower progress with general methods for the conceptual modeling of temporal and spatio-temporal phenomena.

As an example, consider the temporal extended entity relationship (TEER) model [10]. TEER does not add new syntactical constructs to the extended entity relationship model [9] but extends the meaning of existing constructs by adding a temporal dimension. In the TEER model, each entity has associated with it a temporal element that gives the lifespan of the entity. This lifespan can be a single time interval or a union of disjoint time intervals. Each attribute of an entity is then allowed to vary functionally over the lifespan of the entity (i.e., each attribute has associated with it a function from the entity's lifespan to a range of values). In a similar way, relationships have lifespans associated with them, and each relationship attribute can vary temporally over the relationship lifespan, constrained by the lifespans of the participating entities. Also, subclasses have

a constrained temporal relationship with their superclasses. In summary, TEER extends EER by allowing attributes to act as *temporal fluents*, but there is no notation for expressing processes or events.

Another approach to modeling spatio-temporal information is an extension of the object model provided in the TRIPOD project [13] to include the construction of object histories. TRIPOD provides both a model and an implementation for a spatiotemporal database. The underlying data model consists of a collection of temporal snapshots, and takes as primitive the notion of *history*, which models the change of an object, its attributes (also objects), or the relationships that it participates in. Most types can be extended to histories, which provide a sequence of instant or interval timestamped values of the type, thus showing the evolution through time of the instances of the object, attribute, or relationship. TRIPOD moves beyond TEER by providing not just a conceptual model, but also a logical model and implementation, but again there is no notation for explicitly expressing processes or events.

Previous models therefore take objects as the fundamental components of the system, and model change by considering time-varying properties of these objects. For example, a transportation system would be constructed around roads and their time-varying attributes, such as traffic flow rates. The limitation of these approaches is the absence of explicit ways of talking about events and processes, except in so far as events and process affect attributes of objects. The only models which approach the modeling of events in a spatio-temporal system are concerned with the creation, mutation, and destruction of objects [16], and so still have an object bias.

The research question addressed by this paper is how to provide for processes and events in conceptual models of geo-phenomena. This question is addressed in this paper with reference to *dynamic geo-networks*, that is, networks embedded in a (2-dimensional) geographic space. We have chosen to work with networks because although embedded in 2-space and exhibiting many of the properties and behavior of fully spatial structures, they are essentially one-dimensional, and therefore have a simpler structure than more general 2-space entities.

The kinds of application domains we are abstracting from include transportation, utilities, and communication networks. In all these domains, the underlying framework is a network along whose links there are flows. We assume that at a given time, within each link the flow is homogeneous, although we will extend this later to account for “seepage” along a link. In each such application domain, specific theories are well advanced, for example, in the field of transportation, theories seek to describe interactions between the mobile components, e.g., vehicles, vehicle operators, and pedestrians, and the static infrastructure, the highway and its control devices, such as signage, lights, and markings (for example, see [11]). However, what is currently missing is a generic approach to the modeling of the dynamic aspects of geo-networks that provide means of representing not only objects and their temporally varying properties, but also the processes and events in which they participate. The definitions of objects, processes, and events and relationships between these entities provides a rich framework upon which dynamic applications can be modeled.

3 Processes and Events

It is generally agreed that the key concepts required for modeling dynamic phenomena include *object*, *state*, *process* and *event*, but there is little consensus as to exactly how these should be defined. The view taken in this paper is as follows.

Objects and processes have in common the possibility of undergoing change. Change in an object is the most familiar kind, as when, say, a set of traffic lights show red at one time and green at a later time. But a process can also undergo change: for example, the flow of traffic over a bridge is a process, and this can be fast at some times and slow at others. Thus it makes sense to speak of the current *state* of an object or process.

Events differ very significantly in this respect from both objects and processes: it does not make sense to speak of an event as undergoing change. This is because an event is a completed episode of history, and the properties it has, it possesses timelessly. For example, a road accident occurred at a certain location, at a certain time, involving particular cars which were damaged in particular ways. All of these are properties which the event has without temporal qualification; it does not make sense to say that any of them change, and the same goes for all genuine properties of an event.¹

Formally, objects and processes are distinguished from events by the fact that whereas the attributes of objects and processes include both variable (i.e., time-dependent) and constant (i.e., time-independent) properties, the attributes of events are all time-independent. A time-independent attribute may take one of the forms $P(x)$ or $f(x) = v$; here P is a predicate applying to an entity x , and f is a function whose value for entity x is v . Note that time does not feature as an argument in these expressions. A time-dependent attribute, on the other hand, takes one of the forms $P(x, t)$ and $f(x, t) = v$, in which the holding or not holding of predicate P when applied to x , as well as the value of function f for argument x , may vary according to the time given as the second argument. Some examples to illustrate these ideas are:

Time-independent properties of event *accident37*:

location(accident37) = intersection(king_street, union_street)
Involves(accident37, car45)
time(accident37) = t123
type(accident37) = collision

Time-independent properties of processes:

type(roadworks12) = pipelaying
direction(traffic_flow335) = northeast

¹ By ‘genuine’ property we mean intrinsic as opposed to relational properties. Of course an event can ‘change’ from being the most recent car accident to the second most recent one, but this is a change in the relation of the event to the present: it changes because the present changes, not because the event does.

Time-dependent properties of processes:

```
location(roadworks12, t) = king_street.section15
speed(traffic_flow335, t) = 37mph
phase(traffic_light_cycle85, t) = phase3
```

The general picture we are painting is as follows. Objects, states, and processes belong to the ‘snapshot’ view of the world, the world at one time. We can look at the snapshot and say: here we see such-and-such objects, in such-and-such states, undergoing such-and-such processes. Snapshots are continually renewed as time passes, i.e., the snapshot differs from one moment to the next, and this means that the elements present in a snapshot must be capable of undergoing change. By contrast, events belong to the fixed historical record. This is not renewed in the sense of being replaced by a new record, but only in the sense of being incrementally added to as time passes: as events occur, they are added to the record, but once there they are fixed for all time.²

The events that are added to the record are generated by the processes that exist in the continually evolving snapshot. For example, if a process p comes into existence (meaning that it is present in a snapshot for the first time), then the event ‘process p started’ is added to the record. When the process stops, we get the event ‘process p stopped’, and at that point we can also add to the record the complete episode which consists of the lifetime of p from its inception to its termination, whose attributes include such things as its duration and the magnitude of the resulting change (e.g., if p is a motion process, the distance traveled).

The distinction between the evolving sequence of snapshots and the fixed historical record is reminiscent of the SNAP/SPAN ontologies of Grenon and Smith [12], with the crucial difference that we regard processes as belonging in the snapshots whereas in the SNAP/SPAN framework they are placed in the SPAN ontology. For us, a snapshot is something dynamic in that it incorporates ongoing processes (states of change), whereas a SNAP ontology, as we understand it, only incorporates the static properties of the world at one time.

4 Networks and Flows

In this paper a network is defined as a directed multigraph embedded in a surface, with flows associated with the links. The elements of this definition are explained as follows.

A *directed multigraph* is a set of nodes and links such that each link is associated with an ordered pair of nodes; it is a multigraph because we allow more than one link to be associated with the same ordered pair of nodes. Formally, we have a triple $\langle N, L, nodes \rangle$, where $nodes : L \rightarrow N \times N$ is a

² Do not confuse this with the fact that *our* records may undergo changes, e.g., through the correction of errors. By the ‘fixed’ historical record we mean the ideally correct record of what actually happens.

function associating node-pairs with links (intuitively, $nodes(l) = \langle n_1, n_2 \rangle$) tells us that link l connects node n_1 to node n_2). We shall also use the notation $nodes(l) = \langle innode(l), outnode(l) \rangle$.

A directed multigraph is an abstract entity which is not inherently spatial in nature. To ‘spatialize’ it, we must consider it as embedded in space, which in geographical contexts usually means the two-dimensional space of the earth’s surface or some delimited portion thereof. Let S be the space in question, then the embedding is specified by a function loc which (1) maps each node of the multigraph onto a location in S , and (2) maps each link l of the multigraph onto a curve segment whose endpoints are at $loc(innode(l))$ and $loc(outnode(l))$. When we speak of the length of a link, we mean, of course, the length of this curve segment.

The spatially embedded multigraph is a purely static entity. To make it into a dynamic network, we add to each link one or more *flows*, where a flow is conceived as an ongoing movement along the link from its in-node to its out-node—a process in the sense that this term was used earlier. We do not specify exactly what it is that is moving: it may be discrete objects such as vehicles, or it may be some fluid such as water or blood. We do, however, assume that there is some measure of *amount* for whatever is moving through the network, in terms of which we can speak of its density and rate of flow. We require the possibility of associating more than one flow with a given link, so that, for example, we can distinguish between the flow of trucks, the flow of cars, and the flow of bicycles along a given stretch of road: each of these flows has a different *flow-type*.

The flow of type T associated with link l will be denoted $l.flow(T)$. This is a process possessing time-varying attributes such as:

- *speed*: the speed at which the flow passes through the link (measured in distance/time).
- *density*: the amount of flow per unit distance along the link (measured in quantity/distance — the exact meaning of ‘quantity’ will depend on the nature of the flow, e.g., number of vehicles, volume of water).
- *throughput*: the amount of flow passing through the link in unit time (measured in quantity/time).

The syntax of these attributes is illustrated by $speed(l.flow, t) = s$. Clearly $throughput(l.flow(T), t) = speed(l.flow(T), t) \times density(l.flow(T), t)$.

Here we have assumed that the flow attributes are uniform along the length of a link. In reality this may not be the case for two reasons:

- The speed and density of the flow may vary along the length of the link (e.g., in a stretch of road between two sets of traffic signals, the flow near the ends is typically both slower and denser than in the central portion, even if their product, the throughput, is uniform);
- A link may allow *seepage*, i.e., the loss or gain of flow elements *along* the link (e.g., water leaking from a water pipe is a loss; but cars entering a residential street from private driveways situated along its length represents a gain).

To handle seepage adequately we should replace *throughput* by *input* and *output*, related by

$$\text{output}(l.\text{flow}(T), t) = \text{input}(l.\text{flow}(T), t) + \text{seepage}(l.\text{flow}(T), t).$$

(Here seepage *into* the flow is regarded as positive.)

If we want to retain the idealized assumption of uniformity of flow along a link, and still allow for seepage, this can be done by introducing an extra notional node in the middle of the link which can function as a sink or source: then all the seepage is regarded as concentrated at that node, and manifests itself as the difference between the throughputs of its incoming and outgoing flows (i.e., the two links into which the original link was divided).

A flow along a link will generally have a maximum possible throughput, determined by the capacity of the link for that type of flow and general considerations relating to the nature of the flow elements. For example, in a road, there will be a maximum speed defined by the speed limit (in practice rather higher than this!), and at a given speed, there will be a maximum density determined by the average separation of vehicles traveling at that speed (in practice rather less than the officially sanctioned braking distances). Where there is more than one type of flow along a link, these will be as it were in competition with each other for the capacity afforded by the link; the details of this will in general be complicated, and we defer consideration of this to a later occasion.

Nodes may be classified according to the number of incident links. Exceptionally, we might allow a node with only one link. If there is to be a non-zero flow along the link then the node must act as a sink or a source depending on the direction of flow. If there is flow both towards and away from the node, then there must be two separate links to carry those flows (even if, physically, they occupy the same channel, e.g., a narrow road carrying two-way traffic).

If two or more links are mapped to the same pair of nodes (i.e., $\text{nodes}(l_1) = \text{nodes}(l_2) = \dots$), then we shall call them *parallel* links. An example would be in a road network, where there are two or more lanes in the northbound direction between one junction and the next; each of these lanes would be modelled as a separate link, and those links will all be parallel. In general, there will be seepage between these parallel links. Likewise, if $\text{nodes}(l_1) = \langle n_1, n_2 \rangle$ and $\text{nodes}(l_2) = \langle n_2, n_1 \rangle$ we shall say that l_1 and l_2 are *antiparallel* links. A maximal set of links any two of which are either parallel or antiparallel will be called a *linkage*. A linkage incorporates all the links in either direction between a given pair of nodes.

A node which is incident to exactly two linkages corresponds to a control point in the network, that is, some device or circumstance which has the effect of controlling the flow attributes on the incident links. An example would be the start of a speed-limit zone in a road network. Since the flow attributes on either side of the speed-limit sign are different, we must place a node at that point to maintain the uniform flow condition on individual links.

Nodes incident to three or more linkages correspond to junctions or intersections. Assuming there is no seepage at the node itself, then we can say that the

sum of the outputs of a given type from all its incoming links is equal to the sum of the inputs of that type into all its outgoing links.

If we examine the flow *at a node* in detail, we can see that it can be quite complicated. For example, let the incoming links be l_1, l_2, \dots, l_m , and the outgoing links l'_1, l'_2, \dots, l'_n . Then each pair $\langle l_i, l'_j \rangle$ represents a potential component of the flow, i.e., a flow out of l_i and into l'_j . At a four-way road intersection, for example, with incoming links l_1, \dots, l_4 and outgoing links l'_1, \dots, l'_4 , where l_i and l'_i are antiparallel, then the flow of any given type across the intersection actually has twelve components

$$\langle l_1, l'_2 \rangle, \langle l_1, l'_3 \rangle, \langle l_1, l'_4 \rangle, \langle l_2, l'_1 \rangle, \langle l_2, l'_3 \rangle, \langle l_2, l'_4 \rangle, \\ \langle l_3, l'_1 \rangle, \langle l_3, l'_2 \rangle, \langle l_3, l'_4 \rangle, \langle l_4, l'_1 \rangle, \langle l_4, l'_2 \rangle, \langle l_4, l'_3 \rangle$$

If the traffic is dense, it will not be possible for all these flows to proceed simultaneously without collisions occurring. The purpose of traffic lights is to successively enable various subsets of these twelve flows in turn, in such a way that no flows that cross each other are enabled together.

5 Events in a Dynamic Network

We may distinguish three broad classes of events in networks.

- Changes to the structure of the network itself, for example
 - the introduction of a new link
 - the removal (or permanent closure) of a link
 - the creation of a new node dividing an existing link or linkage into two
 - the creation of a new node unconnected with any existing link
 - the removal of a node and all its incident links
 - the removal of a node consequent on a one-to-one merger of its ingoing links with its outgoing links
- After any of these changes, we are in a sense dealing with a different network. Such changes can be expected, in general, to have an impact on the flows even in those parts of the network which have not been changed.
- Changes which do not affect the structure of the network itself but which may affect the flows in the networks; for example
 - introduction of an obstruction in a link, reducing flow along that link
 - temporary closure of a link
 - removal of an obstruction
 - reopening a temporarily closed link
 - introduction of an obstruction at a node which reduces or blocks one or more flow components through that node
 - Finally, there are *flow events*, the changes that occur in the flows themselves (often as a result of events of one of the previous two kinds).
 - creation of a new flow on a link
 - removal of a flow from a link
 - a flow comes to a halt (i.e., the speed becomes zero)

- a flow ‘dries up’ (i.e., the density becomes zero)
- a flow starts up again from a halted state
- a flow starts up again from a dried-up state

Here we may also think about increases or decreases in flow attributes such as speed, density, and throughput. These should be regarded as processes rather than events; but temporally bounded episodes of such increases or decreases may be singled out as events, for example a sudden decrease resulting from the introduction of an obstacle into a link. In particular, we may be interested in peaks and troughs, i.e., events in which some flow attribute attains a local maximum or minimum value. A peak will be preceded by an episode of increasing value and followed by an episode of decreasing value, and *vice versa* in the case of troughs.

6 Causal Relations Amongst Events, States and Processes

In modeling the evolution of a dynamic system such as one of our networks, it is hard to avoid bringing in the notion of causality. An accident on the road causes an obstruction which causes reduced traffic flow. The accident is an event, the obstruction is an object or a state, and the reduced traffic flow is a state or a process. It seems that causal relations can exist between entities of various different types. In this section we attempt to provide a clear account of these causal relations which will, we hope, be adequate for the purpose of modeling network phenomena.

It should be noted first that ‘causes’ is not by any means the only causal relation we are interested in. Worboys and Hornsby [24] list the following event-event relationships: initiation, perpetuation (or facilitation), hindrance (or blocking), and termination. They also list event-object relationships: creation, sustaining in being, reinforcement, degradation, destruction, splitting and merging. Our intention here is to examine some of these relationships more closely, particularly from the point of view of events, states, and processes in networks, and to derive a more systematic classification of the most important relationships.

Consider the following scenario: at time t_1 , the traffic flow at a particular position P on a road network is *high*; at time t_2 , a little later than t_1 , an accident occurs near P , resulting in an obstruction in the road; at time t_3 , a little later than t_2 , the traffic flow at P is *low*.

A close analysis of this scenario reveals the following events:

- E_1 is the occurrence of the accident.
- E_2 is the road’s becoming obstructed.
- E_3 is the reduction of the flow at P from *high* to *low*.

Each of these events occurs at a definite time; for simplicity we shall assume they all occur at t_2 (this is a matter of choosing the temporal granularity of the representation appropriately). In addition, we can identify the following states:

- S_1 is the state of the road being clear
- S_2 is the state of the traffic flow being high
- S_3 is the state of the road being obstructed
- S_4 is the state of the traffic flow being low

These states are all time-dependent, that is, they will hold at some times and not at others. For example, S_1 holds over the interval (t_1, t_2) whereas S_2 holds over the interval (t_2, t_3) . Note that a state may be the state of an object (e.g., the road) or of a process (e.g., the traffic flow).

What are the causal relations amongst these various states and events? We advocate the following interpretation:

- State S_1 *enables* S_2 . (Note that it would be incorrect to say that it *causes* it: the road being clear cannot itself generate the traffic required for high flow!)
- States S_1 and S_2 are *terminated by* events E_2 and E_3 respectively. This is essentially a matter of definition: a state S is naturally terminated by an event E which is defined as the coming into existence of a state incompatible with S .
- Event E_1 *causes* event E_2 , which in turn causes event E_3 .
- Events E_2 and E_3 *initiate* states S_3 and S_4 respectively. This is again a matter of definition.
- State S_3 *perpetuates* (or *maintains*) state S_4 .

Figure 1 shows these relations in diagrammatic form, using rectangular boxes to represent states, and circles to represent events.

From this we can see that each of the causal relations is specific to particular combinations of types, as follows:

- Event-Event: *causes*
- Event-State: *initiates, terminates*
- State-State: *enables, perpetuates*

In addition, we may introduce

- State-Event: *allows, prevents*

These causal relations, shown diagrammatically in figure 2, may be explained as follows. Here, for simplicity, we assume that all events are *punctual*, that is, they occur at an instant rather than over an interval; the explanations would need to be adjusted to allow for events which take place over an interval, i.e., *durative* events.

- ‘ E_i causes E_j ’ means that E_j occurs as a result of E_i ’s occurring. We do not attempt to define ‘as a result of’—see the voluminous philosophical literature on this subject (for example, [20]).
- ‘ E initiates S ’ means that if E occurs at t then (1) S must hold over an interval beginning at t , and (2) there is an interval ending at t throughout which S does not hold.

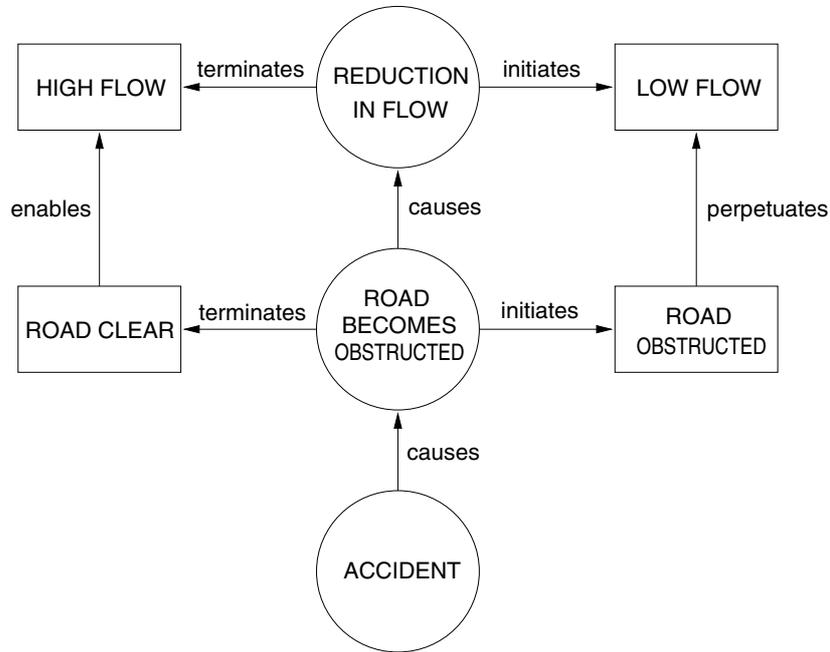


Fig. 1. Diagram of event–state relationships

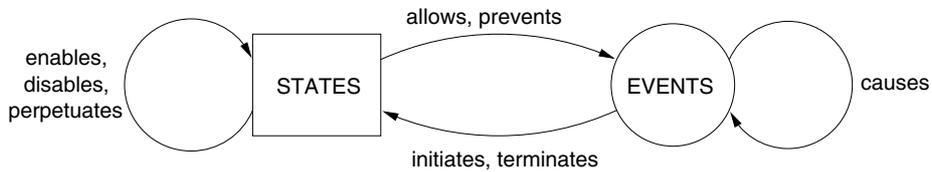


Fig. 2. Kinds of causal relationships between states and events

- ‘ E terminates S ’ means that if E occurs at t then (1) S holds over an interval ending at t , and (2) there is an interval beginning at t throughout which S does not hold.
- ‘ S_i enables S_j ’ means that if S_i holds at t then it is possible for S_j to hold at t .
- ‘ S_i disables S_j ’ means that if S_i holds at t then it is *not* possible for S_j to hold at t .
- ‘ S_i perpetuates S_j ’ means that if S_i and S_j both hold at t , and S_i holds throughout an interval i which starts at t , then S_j will also hold throughout i .
- ‘ S allows E ’ means that if S holds at t then it is possible for E to occur at t .
- ‘ S prevents E ’ means that if S holds at t then it is not possible for E to occur at t .

It should be emphasized that we are not here making hypotheses about the way in which English words such as ‘prevents’ and ‘enables’ are actually used; rather we are stipulating how they *should* be used in the technical context of modeling dynamic networks in an information system. The terms have, however, been chosen so as to accord as nearly as possible to normal English usage.

These type-restrictions may seem to be too narrow: it is easy to come up with examples which are *prima facie* exceptions to them. However, we believe that in every case, a careful analysis will show that our type-restrictions are valid. We illustrate this as follows:

1. Can a state cause an event? Example: the presence of ice on the road caused the car to skid. In fact the presence of ice cannot possibly *cause* the car to skid; the immediate cause must be some maneuver by the driver, e.g., braking. Thus the true picture is that the state of iciness *allowed* the skidding event, and the braking event *caused* it. Of course, we could say that the true cause of the skidding is the conjunction of the ice on the road and the braking; but we believe that our account is simpler and entirely adequate: an event causes another event *in the context of* a state which allows it.
2. Can one event prevent another event? Example: The accident prevented John from arriving at the station on time. This statement leaves out a number of intermediate states and events which are required for a full analysis. The correct picture is as follows:
 - The accident (event) *causes* the road to become obstructed (event)
 - The road’s becoming obstructed (event) *causes* the reduction in the traffic flow (event)
 - The road’s becoming obstructed (event) *initiates* the road’s being obstructed (state)
 - The reduction in the traffic flow (event) *initiates* the low traffic flow (state)
 - The road’s being obstructed (state) *perpetuates* the low traffic flow (state)
 - The low traffic flow (state) *prevents* John’s timely arrival (event)
3. Finally, can a state cause a state? Example: The presence of ice on the road caused the road to be dangerous. Again, a fuller analysis reveals what is going on here:
 - The road’s becoming icy (event)³ *causes* the road to become dangerous (event)
 - The road’s becoming icy (event) *initiates* the road’s being icy (state)

³ ‘The road’s becoming icy’ perhaps refers ambiguously to either an event or a process. Here we mean the event, i.e., the completed transition from a state in which the road is not icy to a state in which the road is icy. This transition may be accomplished by means of a process of incremental accumulation of ice on the road surface; and this process might also be described as the road’s becoming icy.

- The road’s becoming dangerous (event) *initiates* the road’s being dangerous (state)
- The road’s being icy (state) *perpetuates* the road’s being dangerous (state)

The reader may find it helpful to draw diagrams for each of these examples, in a similar manner to Figure 1.

7 Conclusions and Further Work

In this paper we have attempted to identify some of the key general concepts which will be needed to underpin any truly event- and process-oriented model of dynamic geospatial systems. We have concentrated on networks as essentially one-dimensional entities embedded in a two-dimensional space, and have proposed a three-tier representation of these in terms of (1) an abstract graph-like structure, (2) a spatial embedding of this, and (3) a system of flows utilizing the embedded structure.

We then turned our attention to the events which may occur in networks of this kind, carefully distinguishing events, which are delimited episodes in the unchanging historical record, from processes, which are dynamic phenomena existing in the present and subject to change as time passes. We classified network events into changes in the structure of the network itself, temporary disruptions to the flow-carrying capacity of the network, and changes in the flows. An important issue here is how to represent the causal dependencies amongst the various events that can occur and the time-varying states of the objects and processes present in the network. We proposed a set of terms for the various causal relations required here, which we hope may become accepted as a standard terminology in this area.

While we believe that the above work represents a real advance in our understanding of the dynamic aspects of spatial networks, much remains to be done. In particular, the concepts introduced informally in this paper will need to be formalized rigorously in a mathematical (logical or set-theoretic) setting. Then further studies will be needed to show how the general ideas developed here can be applied to specific domains such as traffic (specifically the Integrated Transport Network developed by the Ordnance Survey of Great Britain [18]) or communications systems. Finally, it is evident that some of these ideas will be applicable to wider geospatial domains, beyond our specific focus on networks, and there is much scope for further work in exploring these possibilities.

Acknowledgments

Antony Galton’s work is supported by the Engineering and Physical Sciences Research Council of Great Britain under EPSRC grant number EP/C014502/1. Mike Worboys’ work is supported by the National Science Foundation under NSF grant numbers IIS-0429644, BCS-0327615, and the Ordnance Survey of Great Britain.

References

1. T. Abraham and J. Roddick. Survey of spatio-temporal databases. *GeoInformatica*, 3:61–69, 1999.
2. G. Booch. *Object-Oriented Analysis and Design with Applications*. Addison-Wesley, Reading, MA, 2nd edition, 1993.
3. G. Booch, J. Rumbaugh, and I. Jacobson. *The Unified Modeling Language User Guide*. Addison-Wesley, Reading, MA, 1998.
4. M. Brodie, J. Mylopoulos, and J. Schmidt, editors. *On Conceptual Modeling*. Springer, Berlin, 1984.
5. P. P-S. Chen. The entity-relationship model—toward a unified view of data. *ACM Transactions on Database Systems*, 1(1):9–36, 1976.
6. E. Codd. Extending the relational database model to capture more meaning. *ACM Transactions on Database Systems*, 4(4):397–434, 1979.
7. M. J. Egenhofer and A. Frank. Object-oriented modeling for GIS. *URISA Journal*, 4:3–19, 1992.
8. M. J. Egenhofer and J. R. Herring. A mathematical framework for the definition of topological relationships. In K. Brassel and H. Kishimoto, editors, *Proc. Fourth International Symposium on Spatial Data Handling*, pages 803–813, Zurich, 1990.
9. R. Elmasri and S.B. Navathe. *Fundamentals of Database Systems*. The Benjamin Cummings Publishing Company, 3 edition, 2000.
10. R. Elmasri and G. Wu. A Temporal Model and Query Language for ER databases. In *Proceedings of the Sixth International Conference on Data Engineering*, pages 78–83, 1990.
11. N.H. Gartner, C.J. Messer, and A. Rathi. Traffic flow theory: A state of the art report. Technical report, US Department of Transportation, Federal Highway Administration, www.tfhrc.gov/its/tft/tft.htm.
12. P. Grenon and B. Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.
13. T. Griffiths, A. Fernandes, N. Paton, S.-H. Jeong, N. Djafri, K.T. Mason, B. Huang, and M. Worboys. Tripod: A spatio-historical object database system. In R. Ladner, K. Shaw, and L. Abdelguerfi, editors, *Mining Spatio-Temporal Information Systems*, pages 127–146. Amsterdam: Kluwer, 2002.
14. R. H. Güting, M. Böhlen, C. Jensen, N. Lorentzos, M. Schneider, and M. Vazirgianis. A foundation for representing and querying moving objects. *ACM Transactions on Database Systems*, 25(1):881–901, 2000.
15. R. H. Güting and M. Schneider. Realm-based spatial data types: The ROSE algebra. *VLDB journal*, 4(2):243–286, 1995.
16. K. Hornsby and M.J. Egenhofer. Identity-based change: A foundation for spatio-temporal knowledge representation. *International Journal of Geographical Information Science*, 14(3):207–204, 2000.
17. R. Hull and R. King. Semantic data modeling: Survey, applications, and research issues. *ACM Computer Surveys*, 19:201–260, 1987.
18. Ordnance Survey of Great Britain. *OS MasterMap User Guide*. Romsey Road, Southampton, UK, 2005.
19. J. M. Smith and D. C. P. Smith. Database abstractions: Aggregation and generalization. *ACM Transactions on Database Systems*, 2(2):105–133, 1977.
20. E. Sosa and M. Tooley, editors. *Causation*. Oxford Readings in Philosophy. Oxford University Press, 1993.

21. J.W. van Roessel. Design of a spatial data structure using the relational normal forms. *International Journal of Geographical Information Systems*, 1:33–50, 1987.
22. M.F. Worboys. Object-oriented approaches to geo-referenced information. *International Journal of Geographical Information Systems*, 8(4):385–399, 1994.
23. M.F. Worboys, H. M. Hearnshaw, and D. J. Maguire. Object-oriented data modelling for spatial databases. *International Journal of Geographical Information Systems*, 4(4):369–383, 1990.
24. M.F. Worboys and K. Hornsby. From objects to events: GEM, the geospatial event model. In M. Egenhofer, C. Freksa, and H. Miller, editors, *Third International Conference on GIScience*, volume 3234 of *Lecture Notes in Computer Science*, pages 327–344. Berlin: Springer Verlag, 2004.