PERGAMON

Generating semantics-based trajectories of moving objects

Dieter Pfoser[1]*, Yannis Theodoridis

Computer Technology Institute, PO Box 1122, GR-26221 Patras, Greece

Abstract

The domain of spatiotemporal applications is a treasure trove of new types of data as well as queries. For designing and testing novel data types and access methods that will support these applications, large spatiotemporal datasets are needed. As in many cases it is not possible to obtain real datasets, either they do not exist or they are not accessible, synthetic datasets are governed by artificial parameters rather than by the (sometimes invisible) rules governing real-world behavior. In this work, we show how, at least for some cases, one can generate spatiotemporal datasets that simulate real-world behavior. We illustrate example cases and translate them into appropriate calls of GSTD[2] a spatiotemporal data generator. The generated data is illustrated using two-dimensional snapshot pictures as well as three-dimensional (two spatial plus one temporal dimension) trajectory images.
© 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Data generation; Spatiotemporal datasets; Synthetic data; Moving objects

1. Introduction

Spatiotemporal database management is an emerging technology. Example applications include fleet management, traffic management, and navigational systems, as well as the thriving developments behind mobile computing (Barbará, 1999). The estimates are that by the year 2003, 500 million people will use mobile terminals (Karppinen, 1999). Most of these terminals will be equipped with a GPS device, and, thus, may make their positions available to the outside, digital, and geo-referenced

* Corresponding author.
  E-mail addresses: pfoser@cti.gr (D. Pfoser), ytheod@cti.gr (Y. Theodoridis).
[1] On leave from the Computer Science Department, Aalborg University Denmark.
[2] The GSTD Web interface is found at http://www.cti.gr/rd3.gstd/.

world. Applications, here, include spatiotemporal data mining, as well as providing geo- and time-referenced content.

New techniques comprising data structures and access methods are necessary to handle this novel type of data and queries (Pfoser & Jensen, 2001; Pfoser, Theodoridis, & Jensen, 2000; Theodoridis, Sellis, Papadopoulos, & Manolopoulos, 1998). The testing of these techniques, in turn, requires the existence of datasets, either real or synthetic. Since real datasets (1) are not accessible for some applications and (2) may not be useful for stress testing conditions, a lot of work can be found in the literature on generating synthetic data, following some specifications (desirable cardinality, distribution, etc.).

Work in the area of data generation includes Gray, Sundaresan, Englert, Backlawski, and Weinberger (1994), which aims at generating billion-record SQL databases using C programs running on a shared-nothing computer system consisting of a hundred processors, with a thousand discs. The data generated there is alphanumeric. A generator for Web data, i.e. tables and Web pages, is introduced in Priyadarshini, Qin, Lim, and Ng (1999). Soo (1997) presents a temporal data generator. This work can be seen as the one-dimensional case of spatial data generators such as presented in Günther, Oria, Picouet, Saglio, and Scholl (1998); a web-based spatial data generator that produces sets of rectangles. The generator uses statistical distributions to compute the size, shape, and location of the data. Examples of spatiotemporal data generators are Oporto (Saglio & Moreira, 1999), a network-based generator (Brinkhoff, 2000), and GSTD (Theodoridis, Silva, & Nascimento, 1999). The Oporto generator uses a driving application, the modeling of fishing ships. Ships are attracted by shoals of fish while, at the same time, avoiding storm areas. Fish themselves are attracted by plankton areas. Ships are moving points; shoals, plankton, and storm areas are moving regions. Brinkhoff's network-based generator uses, as its name suggests, an underlying network in which movement occurs. Besides two supplied networks, the user can add its own to the generator. Network edges are of varying weight to model different categories of roads permitting varying movement speeds. This in combination with varying the number and speed of the moving objects allows for a close-to-real-life generation of movement data. The GSTD generator produces moving point or rectangular region objects according to several user-defined parameters. The generator uses various distributions to control the movement and the shape of the objects. Theodoridis et al. (1999) present six scenarios on generating spatiotemporal datasets with respect to the desirable heading of objects, speed, etc. Although useful for access methods testing purposes (Nascimento, Silva, & Theodoridis, 1999), the original algorithm turns out to be limiting with respect to specific real-world behavior, such as directed movement or movement obstructed by buildings or alike objects, which we term *infrastructure*.

In this paper, we extend GSTD to facilitate building such scenarios. In particular, Section 2 provides the background of this work, namely the issue of spatiotemporal objects and the rationale of the GSTD generator, Section 3 proposes appropriate extensions to the algorithm in order to support directed and obstructed movements, while some example scenarios are presented in Section 4. Finally, conclusions and future work are discussed in Section 5.

2. Background

Many different objects in the real world exhibit spatiotemporal characteristics. We can distinguish those objects by, either, their spatial extent or their temporal characteristics. By temporal characteristics we refer to how objects change their properties over time, i.e. continuous or discrete. As for the spatial extent, we distinguish point objects (no spatial extent) and aerial objects (spatial extent). In the real world we can identify point objects but also aerial objects whose extent does not matter in the given application context, and, thus, is reduced to a point; as an application, consider the tracking of the continuous movement of cars, planes, people, etc. Aerial objects, on the other hand, may have an extent that changes over time. Examples of spatially continuously changing objects are storms, troop formations, grasshopper populations, tribes, etc., while examples of spatially discretely changing objects are land parcels (Pfoser & Tryfona, 1998).

An important aspect of a continuous movement is that it cannot be recorded in its entirety. Let us consider the example of a moving object whose position is sampled at points in time, i.e. a moving vehicle or a person, a military unit in the battlefield, etc. In order to record its movement, we would have to know its position at any time, i.e. on a continuous basis. However, GPS and telecommunications technologies only allow us to sample an object's position, i.e. to obtain the position at discrete instances of time, such as every few seconds.

2.1. Trajectories of moving objects

A first approach to represent the movements of objects would be to simply store the position samples. This would mean that we could not answer queries about the objects' movements at times in-between sampled positions. Rather, to obtain the entire movement, interpolation is necessary. The simplest approach is to use linear interpolation, as opposed to other methods such as polynomial splines (Bartels, Beatty, & Barsky, 1987). The sampled positions then become the end points of line segments of polylines, and the movement of an object is represented by an entire polyline in three-dimensional space. In geometrical terms, the movement of an object is termed a *trajectory* (we will use "movement" and "trajectory" interchangeably). The solid line in Fig. 1 represents the movement of a point object. Space and time are combined to form one (three-dimensional) coordinate system. The dashed line shows the projection of the movement on the plane of *x*- and *y*-coordinates (Pfoser & Jensen, 1999). The issue of interpolation becomes more complex when we consider aerial objects. In this case, interpolating instances means to interpolate the shape as well.

Interpolating trajectories raises questions about the uncertainty associated with a particular representation. Did Pfoser and Jensen (1999) describe uncertainty as related to trajectory data, so do Pfoser and Tryfona (2001) address the general issue of uncertainty and imprecision in temporal, spatial, and spatiotemporal contexts.
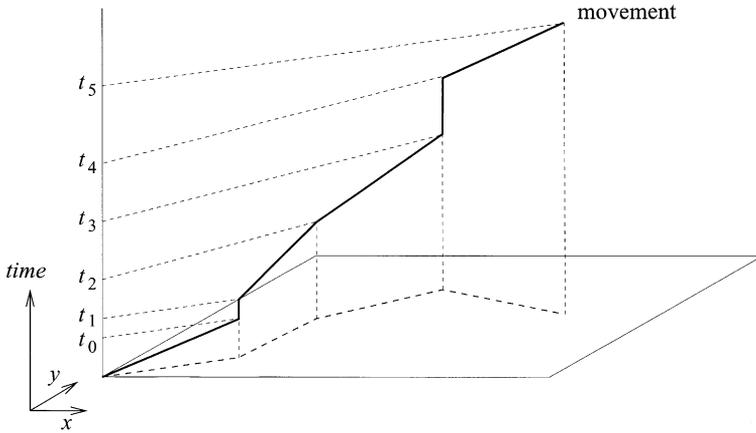
Fig. 1. A moving object trajectory.

2.2. GSTD overview

A fundamental issue in generating synthetic spatiotemporal datasets is the definition of a set of parameters that control the evolution of spatial objects. GSTD (Theodoridis et al., 1999) addresses the following:

- *duration of an object instance*, which involves change of timestamps between consecutive instances,
- *shift of an object*, which involves change of spatial location (in terms of center point shift), and
- *resizing of an object*, which involves change of an object's size (only applicable to aerial objects).

Since the evolution (or "history") of a spatiotemporal object, identified by its *o_id*, is represented by a set of instances (o_id, s_i, t_i), where s_i is the location of object o at instant t_i (s_i and t_i are called *spacestamp* and *timestamp*, respectively), the goal to be reached is the calculation of its consecutive instances, starting from an initial instance (o_id, s_1, t_1). It is also assumed that the spatial workspace of interest is the unit square $[0,1)^2$ and that time varies from 0 to 1 (i.e. the unit interval). For illustration reasons, four instances of a moving object are shown in Fig. 2 together with the corresponding projections on the spatial plane and the temporal axis.

The *shift*, the *duration*, and the *resizing* of an object's instance are represented by the functions:

- **duration**(o_id,interval,curr_tstamp,new_tstamp)
- **shift**(o_id,Δc[ ],curr_sstamp_c,new_sstamp_c)
- **resizing**(o_id,Δext[ ],curr_sstamp_ext[ ],new_sstamp_ext[ ])
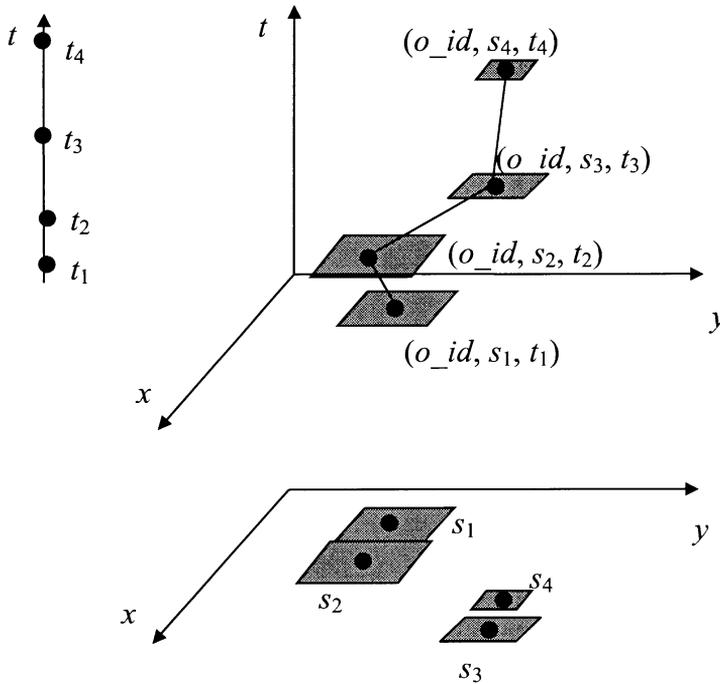
Fig. 2. Consecutive instances of a time-evolving object and the corresponding projections.

which calculate the new timestamp and the new spacestamp's center and extent, called `new_tstamp` (a numeric value), `new_sstamp_c` (a two-dimensional point), and `new_sstamp_ext[ ]` (an array of two intervals), respectively, of an object identified by its `o_id`, as the sums of the respective current values and the respective parameters, namely `interval`, Δc[ ], and Δext[ ].[3]

Not all generated instances of objects are within the workspace. GSTD manipulates *invalid* instances according to one among three alternative approaches: (1) the "*radar*" approach, where coordinates remain unchanged, although falling beyond the workspace, (2) the "*adjustment*" approach, where coordinates are adjusted (according to linear interpolation) to fit the workspace, and (3) the "*toroid*" approach, which assumes a toroidal workspace, as such once an object traverses one edge of the workspace, it enters back in the "opposite" edge.

In summary, GSTD gets several user-defined parameters as input:

- `N` and `D` correspond to the initial cardinality and density (i.e. the ratio of the sum of the areas of data rectangles over the workspace area) of the dataset,

---

[3] The parameters `interval`, `c[ ]`, and `ext[ ]` are as follows: `interval(distr_t(),min_t,max_t)`, `c[ ](distr_c(),min_c,max_c)`, `ext[ ](distr_ext(),min_ext,max_ext)`.

- `starting_id` corresponds to the initial identification number of every object in the dataset,
- `numsnapshots` corresponds to the time resolution of the workspace,
- `distr_init()`, the inital distribution of centers
- `distr_t()`, `min_t`, and `max_t` correspond to the distribution and the domain of the interval parameter,
- `distr_c()`, `min_c[ ]`, and `max_c[ ]` correspond to the distribution and the domain of the Δc[ ] parameter, and
- `distr_ext()`, `min_ext[ ]`, and `max_ext[ ]` correspond to the distribution and the domain of the Δext[ ] parameter.

Several GSTD parameters are controlled by statistical distributions. For example, the initial spatial distribution `distr_init` is a choice among *uniform*, *gaussian*, and *skewed*. For more details on GSTD, please refer to the original paper (Theodoridis et al., 1999) and the GSTD Web interface (Theodoridis & Nascimento, 2000).

3. Introducing semantics in GSTD

Spatiotemporal movement characteristics make it necessary to add new features to the current GSTD implementation. Introducing a *new parameter* called `dinterval` allows us to create more realistic object movements. Additionally, the *infrastructure concept* permits the creation of trajectories stemming from objects moving in an obstructed environment, e.g. cities. We present an infrastructure generator as well as modifications to the GSTD algorithm to consider this information during trajectory generation.

3.1. Extending the GSTD algorithm

The majority of objects in the real world do not move according to statistical parameters but, rather, move intentionally. For example, the movement of a car is guided by the will of its driver, which could be to reach many destinations consecutively. However, the path to such destinations is not linear but according to many factors such as the road network, traffic conditions etc. Overall, however, a movement towards a destination is directed. With the current GSTD implementation such a movement that is overall directed and locally undirected cannot be simulated. Currently, the spatial **shift** parameter, which is responsible for a change of direction, is changed for each new timestamp, i.e. after every `interval`. Thus, by using a broad Δc[ ] parameter the change in **shift** can be larger, and as a consequence, the objects move "nervously". Alternatively, by using a narrow Δc[ ] parameter all objects move in one direction during the whole time frame. To find a middle ground, i.e. objects having directed movements for longer periods of time, we introduce a new parameter called `dinterval`. During a dinterval, Δc[ ] is kept constant. Was the Δc[ ] parameter so far only governed by a distribution function specified as input to GSTD, does it now become a more complex function with the following format:

- Δc[ ](dinterval,new_tstamp,distr_c(),min_c[ ],max_c[ ]

The format of the new parameter is `dinterval(distr_dt(),min_dt,-max_dt)`, thus introducing three extra input parameters, namely `distr_dt()`, `min_dt`, and `max_dt`, corresponding to the distribution and the domain of dinterval.

3.2. Clustered movements

Different groups of moving objects exhibit different kinds of behavior. Let us consider the movements of cars in the morning hours. Parents are dropping their kids at schools in the suburbs, people drive to their jobs downtown, taxis move around in the city center, and people drive to shopping malls. Here, clusters of people exhibit similar movement characteristics.

GSTD indirectly supports clustered datasets. To achieve that, sets of objects moving with varying speed, agility and direction are just to be generated and then appropriately combined into a single dataset.

3.3. Infrastructure

An important assumption in GSTD was that the objects are unobstructed in their movements with the single exception of the workspace boundary. However, in real-world scenarios, there may exist spatial objects that hinder other objects in their movement. We term this set of objects *infrastructure*. Like moving object datasets, infrastructure can be composed by either real or synthetic data.

Focusing on synthetic data, a popular generator for spatial data (used so far in spatial database work) has been implemented. The input parameters of the generator are the following:

- `N`, i.e. the number of rectangles,
- `X` and `Y`, i.e. the (maximum) size of the rectangles in *x*- and *y*- dimension, respectively, and
- `extent_type`, i.e. whether the size of rectangle is constant and equal to `X` or `Y` (`extent_type=c`) or variable up to `X` or `Y` (`extent_type=v`).

The output consists of a set of non-overlapping rectangles with their centers obeying a uniform distribution (other popular distributions, such as gaussian or skewed have been implemented as well).

Fig. 3 depicts an infrastructure file containing 20 rectangles, of a variable box size of up to 20% in both dimensions. The pseudo-code of the infrastructure generator is presented in the Appendix.

When generating moving objects data, each new instance has to be *outside* each of the *infrastructure* objects. To achieve that, we employ a similar approach to what the GSTD algorithm uses in treating objects that are outside the workspace (cf. Section 2.2). From the three originally proposed approaches, it is only the radar and the adjustment approach that can be applied to infrastructure. The application of
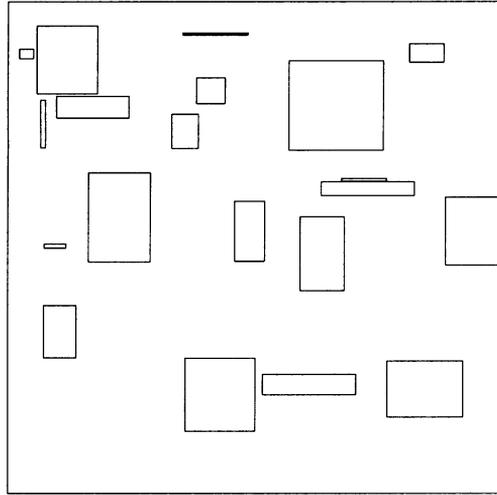
Fig. 3. Infrastructure (a set of rectangles).

the *radar* approach is straightforward; the coordinates of the new instance remain unchanged, although intersecting with infrastructure objects. More important in this context is the "*adjustment*" approach, where coordinates are adjusted (according to linear interpolation) to be outside infrastructure objects.

In this context, do note that we are only dealing with static obstacles, i.e. the infrastructure. Examples of dynamic (or moving) obstacles are chunks of slow moving traffic, motor cades, special transports blocking roads, etc. Those objects could be generated using the GSTD tool itself. However, in such a case the whole environment becomes quite complex.[4]

4. Complex scenarios on moving object trajectories

In this section we show how to generate trajectories of moving objects using GSTD to simulate real-world scenarios. We present various scenarios and depict a proper GSTD parameter setting. The generated data is illustrated using a two-dimensional temporal snapshot representation as well as a three-dimensional polyline representation.

An initial goal of data generators is the control of basic object properties, such as speed, heading, and agility of objects. Subsequently, we combine the groups of objects showing similar movements characteristics to form clusters. The last set of movement examples is characterized by infrastructure hindering the free movement

---

[4] Algorithms determining whether a moving object intersects a moving obstacle fall in the category of a *spatiotemporal join*, which deserves careful investigation (Theodoridis et al., 1998).

of objects. We use a random set of rectangles to represent static real-world objects, e.g. the movements of cars would be obstructed by buildings, rivers, parks, etc.

For all the examples to follow, we have made some assumptions: (1) The GSTD generator can generate both, point objects and aerial objects. The extent of an object is of no importance in the present context, thus in the sequel we will only deal with *point objects*. (2) The GSTD generator employs several approaches of how to adjust the positions of objects leaving the workspace. As already mentioned, we have chosen to experiment with the *adjustment* approach. (3) We assume that the spatial workspace of interest is the unit square $[0,1)^2$ and that time varies from 0 to 1 (i.e. the unit interval). Finally, (4) the format in which the distributions are specified in the following examples is identical to the actual input stream used in GSTD, i.e. `r` stands for a *uniform* distribution, `g_0.5_0.1` stands for a *gaussion* distribution with a mean = 0.5 and a standard deviation = 0.1, and `s(1)` stands for a *skewed* distribution.

4.1. Movement parameters

The first set of examples demonstrates how basic characteristics of objects such as speed, agility, and direction can be manipulated using the parameters of the GSTD generator.

4.1.1. Varying objects speed

Objects do not move with constant speed, e.g. cars are parked, stopped at traffic lights, move in cities, or on highways. Further, different kinds of objects have different speed capabilities. Cyclists cannot move as fast as cars, different kinds of ships move at different speeds, e.g. high-speed ferries and fishing boats. The scope of this section is to illustrate the connection between the real-world *speed* characteristic and the relevant GSTD parameters.

A varying object speed can be achieved by varying the `distr_c()`, `min_c[ ]`, and `max_c[ ]` parameters in GSTD. This means that during a temporal interval the center of an object can move by as much as the distribution permits. The distribution can be any one supported by GSTD. In the example shown in Fig. 4(a), we chose `distr_c()` to be uniform, with `min_c[x,y]=(-0.3, -0.3)` and `max_c[x,y]=(0.3, 0.3)`. However, do note that these parameters have to be seen in connection with the given temporal interval determined by `distr_t()`, `min_t`, and `max_t`.

The snapshots of Fig. 4(a) are taken at timestamps 0 (initial distribution), 0.25, 0.50, 0.75, and 1 (final distribution). Fig. 4(b) gives a 3D representation of the trajectories. The two spatial dimensions are in the horizontal plane, whereas the temporal dimension is in the vertical direction. The cube represents the extents of the workspace. To give a better visual impression of the movement, the spacestamps are connected using linear interpolation to form polylines.

Fig. 5 illustrates faster moving objects. As before, we chose `distr_c()` to be uniform, but with `min_c[x,y]=(-0.8, -0.8)` and `max_c[x,y]=(0.8, 0.8)`. The possible range of deviation is larger. One can see that after an initial dispersion phase, the distribution of objects is almost uniform throughout the workspace. This is in contrast to the previous example for slower movement, where

(a)


```
2D: N = 100; 3D: N = 20
D = 0
starting_id = 1
numsnapshots = 1000
distr_init()= g_0.5_0.01

Distr_t() =       distr_dt() =
g_0_0.03          g_0_0
min_t = 0         min_dt =0
max_t = 1         max_dt = 0

Distr_c() = r     distr_ext() = r
min_c[] =         min_ext[] = 0,0
-0.03,-0.03       max_ext[] = 0,0
max_c[] =
0.03,0.03
```

(b)   (c)
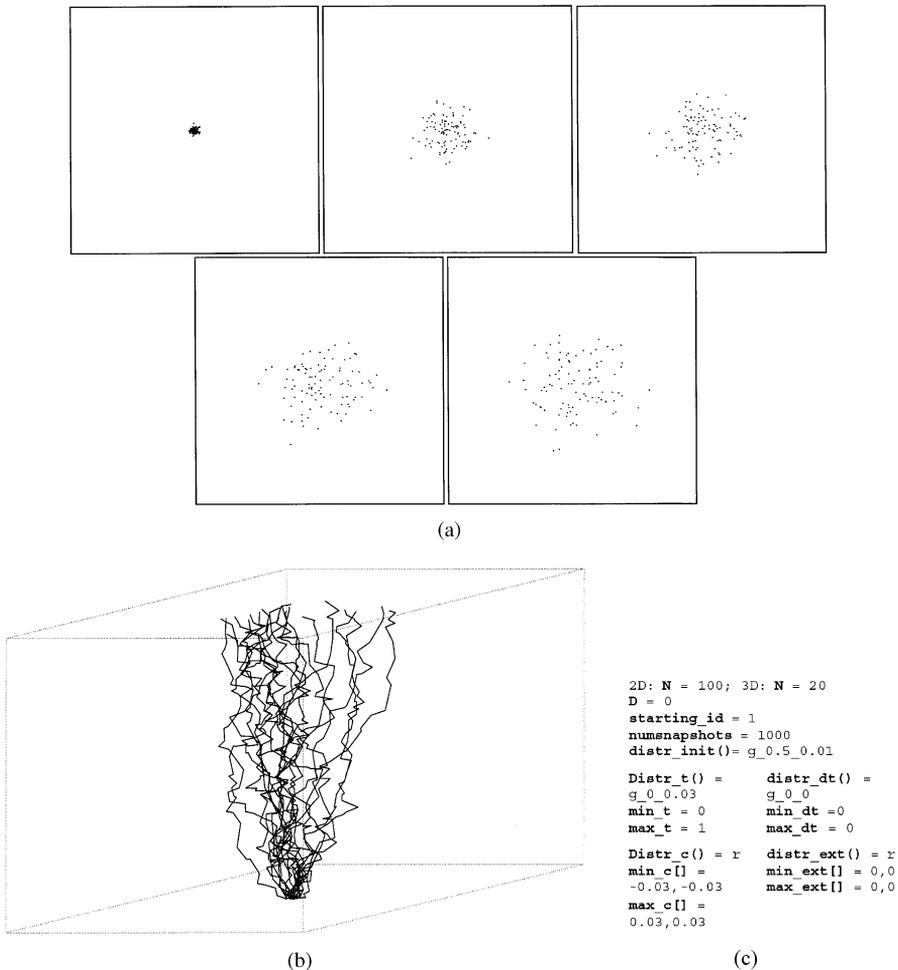
Fig. 4. Slow moving objects. (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

the initial distribution highly affects the overall positioning of the objects. Fig. 5(b) shows a rather nervous movement of objects, i.e. a frequent change of the heading. With the next set of examples we show how to smoothen directional movement.

4.1.2. Agility of objects

To generate trajectories for more directed movements, an additional GSTD parameter, distr_dt(), was introduced in Section 3.1. For the example shown in Fig. 6, we used distr_dt()=g_0.2_0.1, i.e. a gaussian distribution with a mean = 0.2 and a standard deviation = 0.1, with min_dt=0 and max_dt=1. In general, the larger a dinterval, the longer an object moves into a specified direction, i.e. with the average dinterval=0.2, there are five changes of the general

(a)


```
2D: N = 100; 3D: N = 20
D = 0
starting_id = 1
numsnapshots = 1000
distr_init() = g_0.5_0.01

distr_t() =            distr_dt() =
g_0_0.03              g_0_0
min_t = 0            min_dt =0
max_t = 1            max_dt = 0

distr_c() = r         distr_ext() = r
min_c[] =            min_ext[] = 0,0
-0.08,-0.08          max_ext[] = 0,0
max_c[] =
0.08,0.08
```

(b)   (c)
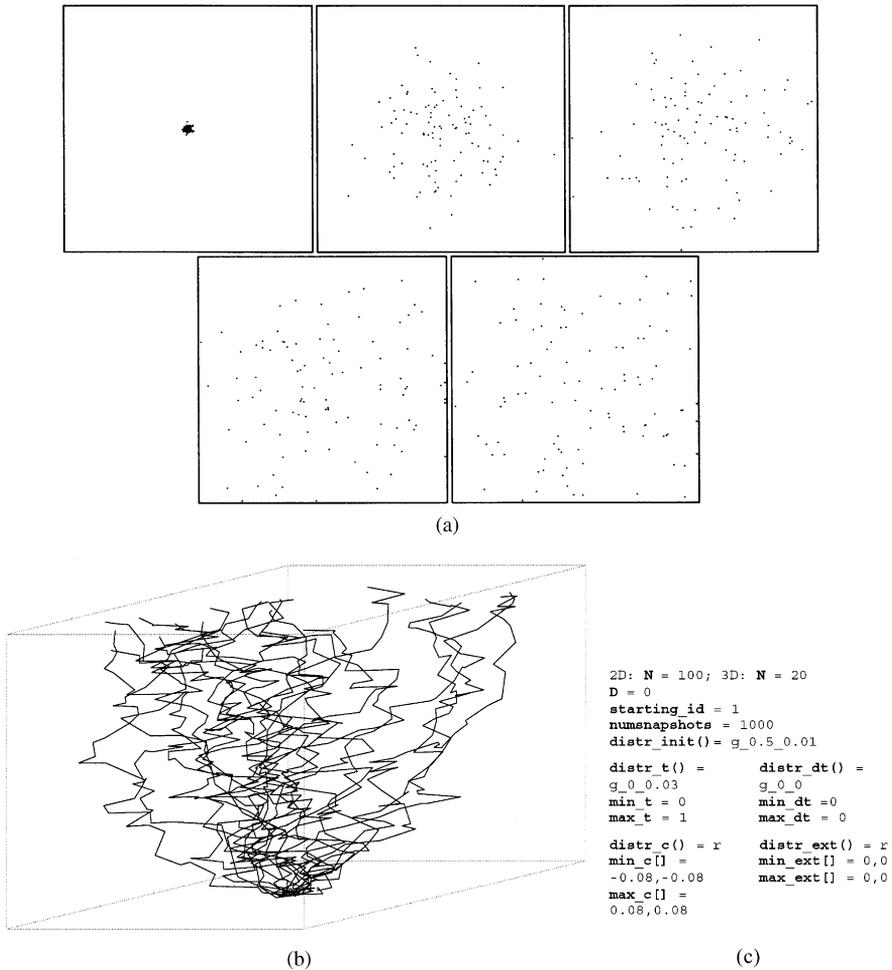
Fig. 5. Fast moving objects. (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

direction during the whole time frame. The other parameters are the same as in the example of Fig. 4, depicting slow movement.

The overall effect of the `dinterval` parameter is that a chosen direction is kept for a larger time interval. Thus, in comparing Figs. 4 and 6, one can easily see that the trajectories now appear to be less "nervous," and, at the same time, the objects move away from their initial positions. In contrast to this, the objects in Fig. 4 move around their inital positions; the many changes of Δext[ ] , with a mean = 0, cancel each other.

4.1.3. Preferred direction

Direction is an important characteristic of a movement. To achieve directional movement we, again, manipulate the `min_c[ ]` and `max_c[ ]` parameters. Conse-

(a)


```
2D: N = 100; 3D: N = 20
D = 0
starting_id = 1
numsnapshots = 1000
distr_init()= g_0.5_0.01

distr_t()=          distr_dt()=
g_0_0.03            g_0.2_0.1
min_t = 0           min_dt = 0
max_t = 1           max_dt = 1

distr_c()= r        distr_ext()= r
min_c[] =           min_ext[] =
0.03,-0.03          0,0
max_c[] =           max_ext[] =
0.03,0.03           0,0
```

(b)   (c)
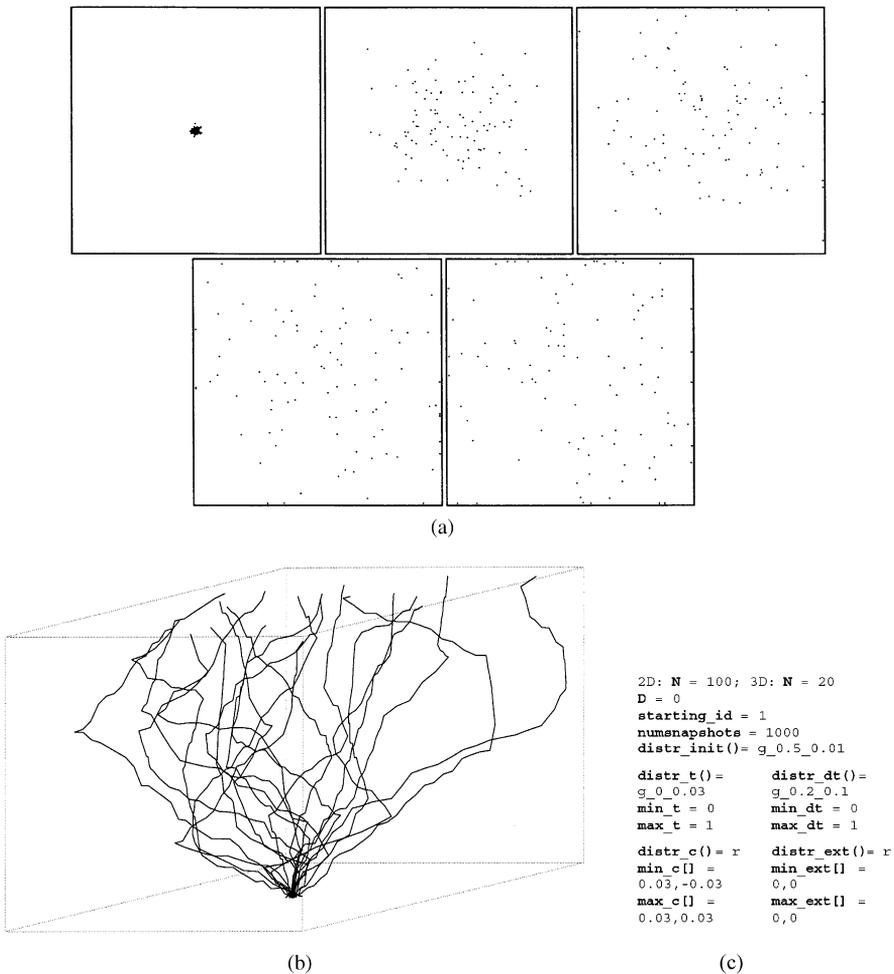
Fig. 6. Smoothly moving objects. (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

quently, we can adjust four parameters, `-x, x, -y`, and `y`, to bias movement. In Fig. 7, a directed movement is illustrated. The parameter setting is `distr_c()` to be uniform and `min_c[x,y]`=(−0.03, −0.1) and `max_c[x,y]`=(0.03, 0.03), i.e. the distribution in the y direction is biased towards the negative direction. General directional movements can be achieved by combining smaller/larger deviation values.

4.1.4. Summary

The objective of this section was to show the correspondence between movement characteristics and GSTD parameters (cf. Table 1).

The speed of objects is determined by the fraction $\Delta c[\ ]$ `/interval`. Thus, in the examples, we specify ranges for both parameters, i.e. distributions. Slower and faster

(a)


```
2D: N = 100; 3D: N = 20
D = 0
starting_id = 1
numsnapshots = 1000
distr_init() = g_0.5_0.01

distr_t()=        distr_dt()=
g_0_0.03          g_0_0
min_t = 0         min_dt = 0
max_t = 1         max_dt = 0

distr_c() = r     distr_ext() = r
min_c[] =         min_ext[] = 0,0
-0.03,-0.1        max_ext[] = 0,0
max_c[] =
0.03,0.03
```

(b)   (c)
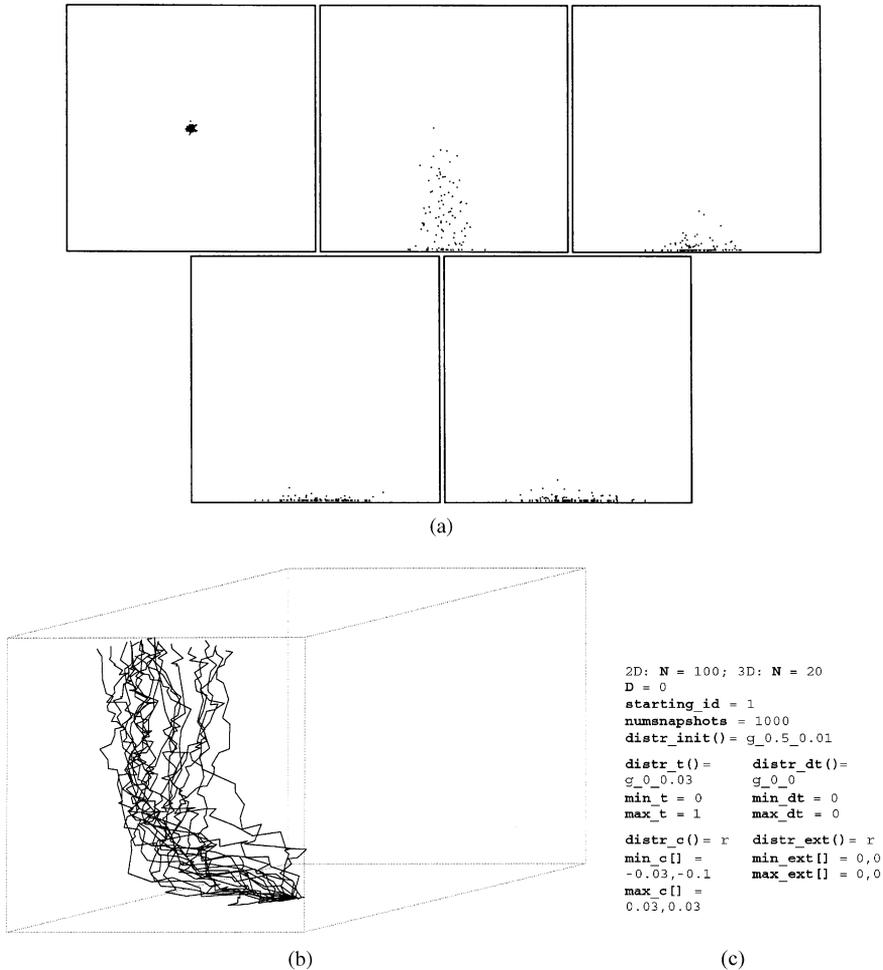
Fig. 7. A directed movement towards "south". (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

moving objects are characterized by allowing a smaller/larger range for $\Delta c[\ ]$ by at the same time keeping interval constant.

The agility of objects is governed by the new dinterval parameter; the larger the range, the fewer times the objects change their general direction. The immediate direction is governed by the $\Delta c[\ ]$ parameter; its values determine the direction, e.g. if $\Delta c[\ x] = 0.03$ and $\Delta c[\ y] = 0.03$, the object moves into a "Northeastern" direction.

4.2. Clustered movement

In the previous section we assumed that the movement of objects within one example is determined by one set of parameters. In this section, we show how to use

Table 1
GSTD parameters and their influence on movement characteristics

Parameter/Characteristic	Speed	Agility	Direction
`interval: distr_t(), min_t, max_t`	√		
`dinterval: distr_dt(), min_dt, max_dt`		√	
`Δc[]: distr_c(), min_c[], max_c[]`	√		√

Table 2
Infrastructure parameters

Example	Parameters	Illustration
Sparse infrastructure	N = 20 X = 0.2 Y = 0.2 `extent_type=v`	
Dense infrastructure	N = 40 X = 0.2 Y = 0.2 `extent_type=v`	
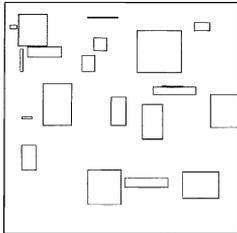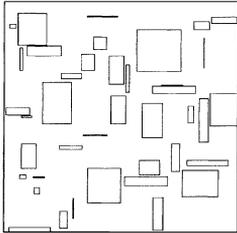
the capabilities of GSTD to generate clustered movement of objects within one dataset.

The key parameter for this purpose is called `starting_id`. By generating various datasets with different movement characteristics and appropriate `starting_id` values, we can then use a simple external concatenation program to combine them. Fig. 8 visualizes the so generated dataset of four clusters of objects moving into four different directions. In Fig. 8(b) one can observe that the objects "crawl" up the corners. The movement parameters, listed in Fig. 8(c), direct the objects to the corners. Because of the adjustment approach they are forced "inside" when they try to leave the workspace, but try again to leave because of the parameter setting.

4.3. Obstructed movement

In Sections 4.1 and 4.2, we assumed that the objects could move freely in the workspace. In this section, we illustrate how infrastructure can hinder the objects in

(a)

(b)

```
2D: N = 200; 3D: N = 20
D = 0
starting_id = 1/6/11/16
numsnapshots = 1000
distr_init()= g_0.5_0.01

distr_t()=            distr_dt()=
g_0_0.03             g_0_0
min_t = 0            min_dt = 0
max_t = 1            max_dt = 0

Distr_ext()= r
min_ext[] = 0,0
max_ext[] = 0,0

distr_c()= r
4 clusters:
1st min_c[] = -0.065,-0.065
 max_c[] = 0.03,0.03
2nd min_c[] = -0.03,-0.065
 max_c[] = 0.065,0.03
3rd min_c[] = -0.03,-0.03
 max_c[] = 0.065,0.065
4th min_c[] = -0.065,-0.03
 max_c[] = 0.03,0.065
```

(c)
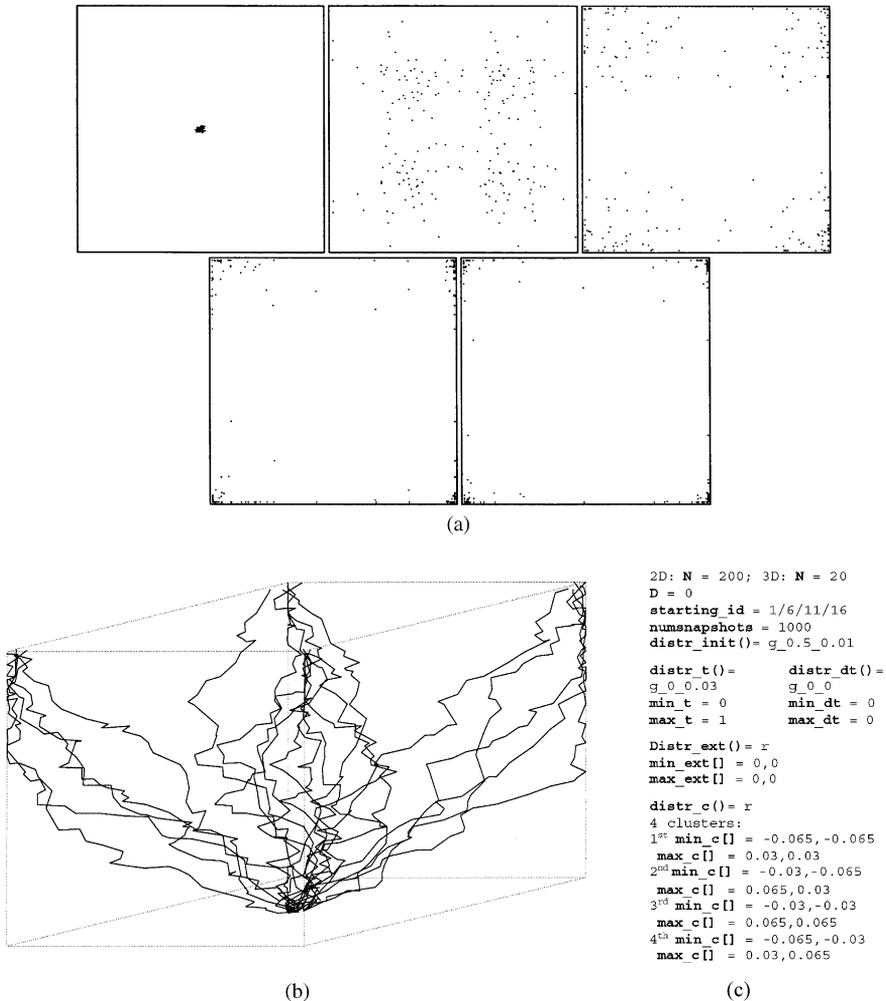
Fig. 8. Clustered movements towards the corners. (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

their movement (cf. Section 3.3). In particular, Table 2 lists the parameters to generate two infrastructure datasets containing 20 and 40 objects, respectively. The size of the objects is limited to 20% of the size of the workspace per dimension.

The first example represents sparse infrastructure. We generated the datasets visualized in Fig. 9. In this example, we chose the "adjustment" approach, i.e. when GSTD computes a new position that is inside an infrastructure object, "adjustment" moves it "outside", where "outside" in this context means onto the border of the obstacle [cf. the first snapshot in Fig. 9(a)].

The movement parameters in this example are the same as in Section 4.1. In comparing Fig. 9 with Fig. 6, we can observe that the infrastructure deflects the

(a)

(b)

2D: **N** = 100; 3D: **N** = 20
D = 0
starting_id = 1
numsnapshots = 1000
distr_init() = g_0.5_0.01

distr_t() =   **distr_dt()** =
g_0_0.03   g_0.2_0.1
min_t = 0   **min_dt** = 0
max_t = 1   **max_dt** = 1

distr_c() = r   **distr_ext()** = r
min_c[] =   **min_ext[]** = 0,0
0.03,-0.03   **max_ext[]** = 0,0
max_c[] =
0.03,0.03

(c)
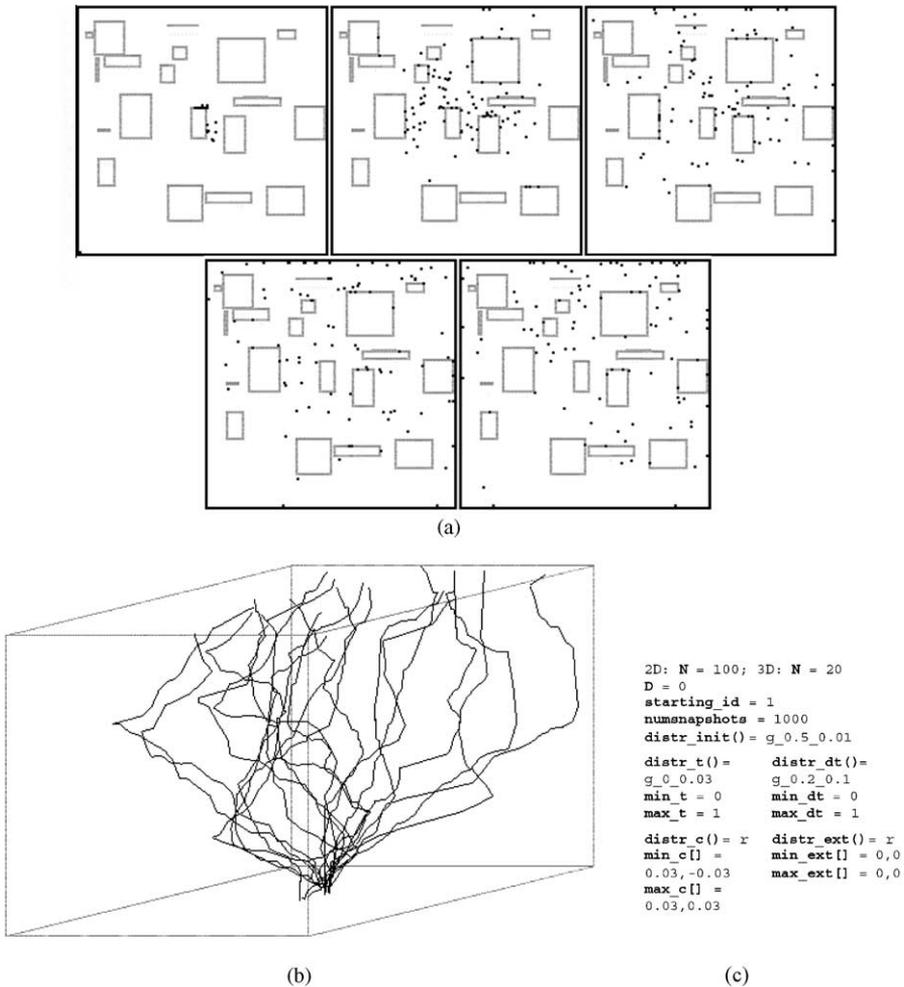
Fig. 9. Obstructed movement, sparse infrastructure. (a), 2D snapshots; (b), 3D trajectories; (c) GSTD parameters.

objects towards the upper right corner. Also, by comparing Fig. 9(b) and Fig. 6(b), we conclude that although the movements are similar (the trajectories resemble each other), the trajectories influenced by infrastructure are "edgier."

In the visualization of the trajectories in Fig. 9(b) we use linear interpolation to connect the spacestamps, i.e. although it is guaranteed that all the spacestamps lie outside the infrastructure objects, this does not necessarily have to be the case for the interpolating line segments.

Fig. 10 visualizes a dataset with the same movement parameters as before, but using the dense infrastructure dataset containing 40 objects (cf. Table 2).
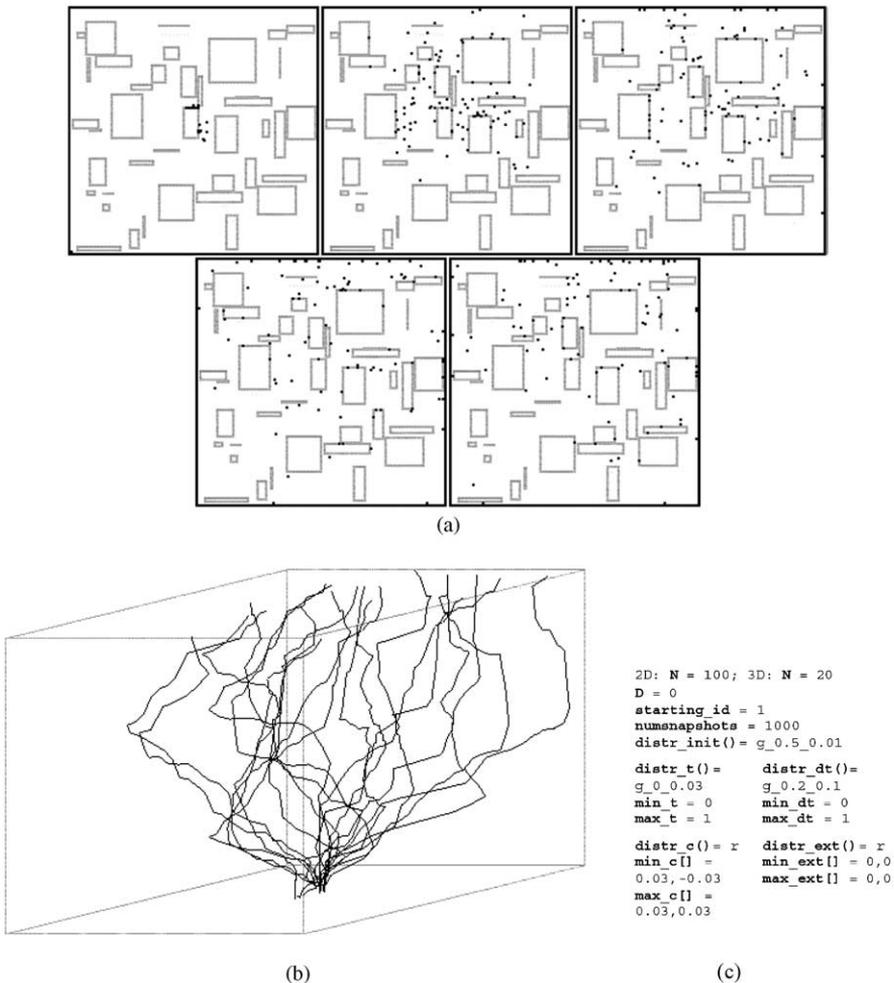
Fig. 10. Obstructed movement, dense infrastructure. (a), 2D snapshots; (b), 3D trajectories; (c), GSTD parameters.

5. Conclusions and future work

This work presents several examples of how real-world movement characteristics can be translated into appropriate parameters of a spatiotemporal data generator. In this work, we used the GSTD tool to generate trajectories of moving point objects. To accommodate different degrees of object agility, we introduced an additional GSTD parameter, called `dinterval`, over the original implementation. Furthermore, we illustrated how clustered movement can be supported. Finally, we noticed that object movement in the real world is not free of

obstacles, here termed *infrastructure*. Thus, as a third extension, we presented an appropriate modification to the GSTD tool to accommodate infrastructure during data generation.

A major part of the work is dedicated to showing different movement examples, such as varying speed, agility, direction, as well as clustered and obstructed movement. The various examples are visualized using a two-dimensional snapshot plus a three-dimensional polyline representation of the trajectories.

This work points to several future research directions. An immediate goal is to extend the GSTD tool provided through a Web interface (Theodoridis & Nascimento, 2000) to accommodate the presented modifications.

Spatiotemporal data generators are parameterized using mathematical, better, statistical concepts. On the other hand, several real-world movements are governed by human behavior. An interesting approach would be to parameterize these behavioral patterns as they can be used for data generation. A somewhat related approach to structuring space can be found in (Raubal, Egenhofer, Pfoser, & Tryfona, 1998). Particular semantic concepts that could be studied in this context are *closeness*, e.g. objects that lie close to each other or close to an obstacle are treated differently, and *attraction*, e.g. objects change direction when they approach a target. Since these concepts raise neighborhood issues, work in nearest-neighbor (Roussopoulos, Kelley, & Vincent, 1995) as well as closest-pair queries (Corral, Manolpoulos, Theodoridism, & Vassilakopoulos, 2000) is worth to be integrated into GSTD.

As we outlined in Section 2.1, one might use linear interpolation to estimate the position in between known instances. Consequently, one cannot assume that the "real" trajectory of a moving object coincides with the polyline interpolation. In other words, every new instance is adjusted with respect to the previous instance and the underlying infrastructure. It could be further adjusted such that even the interpolating line segment does not intersect with any infrastructure object. The adjustment algorithm in the current GSTD implementation does not consider this kind of semantics, this is a task for near future work. Also, the current GSTD implementation uses only static infrastructure. As outlined in Section 3.3, infrastructure can be dynamic and extending the GSTD tool towards this direction would be a logical step.

Overall, data generation initiated in (Theodoridis et al., 1999) is just but one tool in the larger framework of benchmarking access methods and query processing techniques. In the area of spatiotemporal databases, this effort started within the ChoroChronos project (Frank et al., 1999) and is still in progress (Nascimento et al., 1999; Papadopoulos, Rigaux, & Scholl, 1999).

Acknowledgements

Appendix

The infrastructure generator

```
Generate_Infrastructure algorithm
Input:    values N, flag constant
          arrays extent[]
Output:   instance(low[], high[])

begin
01  for N do
02   Set low[] = random(0,1)
03   if constant = 1 then //constant size
04   Set high[] = low[] + extent[]
05   else //variable size
06   Set delta[] = random(extent[])
07   Set high[] = low[] + delta[]
08   end-if
09   if low >= 0 and high <= 1 then //within workspace
10   output(low[],high[])
11   end-if
12  end-for
end.
```

The modified GSTD algorithm (Theodoridis et al., 1999)

```
Generate_Spatio_Temporal_Data algorithm
Input:    values N, starting_id, numsnapshots, D,
          min_t, max_t, min_dt,max_dt
          arrays min_c[], max_c[], min_ext[], max_ext[]
          distributions distr_init, distr_t, distr_dt, distr_c,
distr_ext
Output:   instance (id, t, l_l_point, u_r_point), validity_flag

begin
01 for each id in range [starting_id .. N+starting_id] do
02  /*initialization phase*/
03  Set c[] = RNG(distr_init(), 0, 1), ext[] = extent(N, D)
04  Set t = 0, active = TRUE
05 end-for
06 Set step = 1 / numsnapshots
07 for each id in range [starting_id .. N+starting_id] do
08  /*loop phase*/
09  Set next_snapshot = step
10  while active do
11  /* calculate delta-values and new instances */
12  Set interval = RNG(distr_t(), min_t, max_t)
13  if ∑ interval > dinterval
14  Set dinterval = RNG(distr_dt(), min_dt, max_dt)
15  Set Δc[] = RNG(distr_c(), min_c[], max_c[])
16  Set Δext[] = RNG(distr_ext(), min_ext[], max_ext[])
17  reset ∑ interval
18  end-if
19  Set old_instance = instance
20  update_instance(instance)
21  /* check instances and output */
22  if t > 1 then
23  active = FALSE
24  output(old_instance, current[i], next_snapshot)
25  else /*check instance validity and output*/
26  Set validity_flag = valid(instance)
27  if validity_flag = FALSE and approach ≠ "radar" then
28  adjust_coords(instance, approach)
29  end-if
30  if t > next_snapshot then
31  output(old_instance, current[i], next_snapshot)
32  end-if
33  end-if
34  end-while
35 end-for
end.
```

References

Barbará, D. (1999). Mobile computing and databases—a survey. *IEEE Transactions of Knowledge and Data Engineering, 11*(1), 108–117.

Bartels, R., Beatty, J., & Barsky, B. (1987). *An introduction to splines for use in computer graphics and geometric modeling*. Los Altos, California: Morgan Kaufmann Publishers.

Brinkhoff, T. (2000). Generating network-based moving objects. In *Proceedings of the 12th International Conference on Scientific and Statistical Database Management* (pp. 253–255).

Corral, A., Manolopoulos, Y., Theodoridis, Y., & Vassilakopoulos, M. (2000). Closest pair queries in spatial databases. In *Proceedings of the ACM-SIGMOD Conference on Management of Data* (pp. 189–200).

Frank, A., Grumbach, S., Güting, R., Jensen, C., Koubarakis, M., Lorentzos, N, Manolopoulos, Y., Nardelli, E., Pernici, B., Schek, H.-J., Scholl, M., Sellis, T., Theodoulidis, B., & Widmayer, P. (1999). Chorochronos: a research network for spatiotemporal database systems. *SIGMOD Record*, *28*(1), 110–114.

Gray, J., Sundaresan, P., Englert, S., Backlawski, K., & Weinberger, P. (1994). Quickly generating billion-record synthetic databases. In *Proceedings of the ACM-SIGMOD Conference on Management of Data* (pp. 243–252).

Günther, O., Oria, V., Picouet, P., Saglio, J.-M., & Scholl, M. (1998). Benchmarking spatial joins *À La Carte*. In *Proceedings of the 10th International Conference on Scientific and Statistical Database Management* (pp. 32–41).

Karppinen, J. (1999). Wireless multimedia communications: a Nokia view. In *Proceedings of the Wireless Information Multimedia Communications Symposium*, Aalborg University.

Nascimento, M., Silva, R., & Theodoridis, Y. (1999). Evaluation of access structures for discretely moving points. In *Proceedings of the International Workshop on Spatio-Temporal Database Management* (pp. 171–188).

Papadopoulos, A., Rigaux, P., & Scholl, M. (1999). A performance evaluation of spatial join processing strategies. In *Proceedings of the 6th International Symposium on Spatial Databases* (pp. 286–307).

Pfoser, D., & Jensen, C. (1999) Capturing the uncertainty of moving-object representations. In *Proceedings of the 6th International Symposium on Spatial Databases*. (pp. 111–132).

Pfoser, D., & Jensen, C. (2001). Querying the trajectories of on-line mobile objects. In *Proceedings of the 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access* (pp. 66–73).

Pfoser, D., Theodoridis, Y., & Jensen, C. (2000). Novel approaches to the indexing of moving object trajectories. In *Proceedings of the 26th International Conference on Very Large Databases*, (pp. 395–406).

Pfoser, D., & Tryfona, N. (1998). Requirements, definitions, and notations for spatiotemporal application environments. In *Proceedings of the 6th International Symposium on Advances in Geographic Information Systems* (pp. 124–130).

Pfoser, D., & Tryfona, N. (2001). Capturing fuzziness and uncertainty of spatiotemporal objects. In *Proceedings of the 5th East European Conference on Advances in Databases and Information Systems* (pp. 112–126).

Priyadarshini, P., Qin, F., Lim, E., & Ng, W. (1999). WEDAGEN—a synthetic Web database generator. In *Proceedings of the 10th International Workshop on Database & Expert Systems Applications* (pp. 744–748).

Raubal, M., Egenhofer, M., Pfoser, D., & Tryfona, N. (1998). Structuring space with image schemata: wayfinding in airports as a case study. In *Proceedings of the International Conference on Spatial Information Theory* (pp. 85–102).

Roussopoulos, N., Kelley, S., & Vincent, F. (1995). Nearest neighbor queries. In *Proceedings of the 1995 ACM SIGMOD Conference on Management of Data* (pp. 71–79).

Saglio, J.-M., & Moreira, J. (1999). Oporto: a realistic scenario generator for moving objects. In *Proceedings of the 10th International Workshop on Database & Expert Systems Applications* (pp. 426–432).

Soo, M. (1997). *Generating temporal data: a white paper*. University of South Florida: Department of Computer Science and Engineering.

Theodoridis, Y., & Nascimento, M. (2000). Generating spatiotemporal datasets on the WWW. *SIGMOD Record*, *29*(3), 39–43.

Theodoridis, Y., Sellis, T., Papadopoulos, A., & Manolopoulos, Y. (1998). Specifications for efficient indexing in spatiotemporal databases. In *Proceedings of the 10th International Conference on Scientific and Statistical Database Management* (pp. 123–132).

Theodoridis, Y., Silva, R., & Nascimento, M. (1999). On the generation of spatiotemporal datasets. In *Proceedings of the 6th International Symposium on Spatial Databases* (pp. 147–164).