

Situational Visualization

David M. Krum, William Ribarsky, Christopher D. Shaw*, Larry F. Hodges, and Nickolas Faust
GVU Center, Georgia Institute of Technology,
801 Atlantic Drive, Atlanta, Georgia, 30332-0280, USA
{dkrum,ribarsky,cdshaw,hodges}@cc.gatech.edu, nick.f Faust@gtri.gatech.edu
Phone: 404-894-6328, Fax: 404-894-2970

ABSTRACT

In this paper, we introduce a new style of visualization called *Situational Visualization*, in which the user of a robust, mobile visualization system uses mobile computing resources to enhance the experience and understanding of the surrounding world. Additionally, a Situational Visualization system allows the user to add to the visualization and any underlying simulation by inputting the user's observations of the phenomena of interest, thus improving the quality of visualization for the user and for any other users that may be connected to the same database. Situational Visualization allows many users to collaborate on a common set of data with real-time acquisition and insertion of data. In this paper, we present a Situational Visualization system we are developing called Mobile VGIS, and present two sample applications of Situational Visualization.

Keywords

Mobile users and collaborators, dynamic databases, synchronized databases, real-time acquisition and insertion of data, location-specific services, location and time-specific user input.

1. INTRODUCTION

Many tasks require a detailed knowledge of the local environment as well as an awareness of rapidly changing and interacting events. This awareness is termed situational awareness or situation awareness. It can be formally defined as:

“the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future.” [5]

This type of spatio-temporal knowledge and awareness is important in many tasks, including aircraft piloting, military

*Contact author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VRST'01, November 15-17, 2001, Banff, Alberta, Canada.
Copyright 2001 ACM 1-58113-427-4/01/0011 ...\$5.00.

maneuvers, law enforcement, fire fighting, on-site construction engineering, and even the daily automobile commute.

In this paper, we define a set of techniques to provide computer-based support for situational awareness in augmented and virtual reality systems. We wish to bring large, diverse sets of data into such systems, and manage the limited visualization resources of the system and the limited perceptual and cognitive resources of the user. We describe this body of techniques as *Situational Visualization*. Situational Visualizations are responsive to a user's location in space and time, as well as her actions and situation. They integrate real-time and static data from a variety of sources and allow users to view time-varying events and objects and to project trends into the future.

The difference between such Situational Visualizations and other types of visualizations (e.g. information visualization and scientific visualization) is that the user is situated in the environment and experiencing the situation firsthand and in real time. These elements of Situational Visualization become especially significant for mobile users who can move from place-to-place and can encounter new and unexpected situations. The user must be able to receive timely information associated with the situations encountered.

From the user's perspective, this implies that the system must be:

Unobtrusive: The user must be able to attend to the environment without being distracted by the system interface.

Context-Aware: The system must be aware of the user's location and circumstances.

Time-Aware: The system must help the user anticipate events, objects, and trends.

Priority-Aware: The system must help the user prioritize and attend to the most important events and objects.

Users can also act as data collectors, having a hand in the construction of the database. It is now possible for users to carry a variety of sensors (cameras, various environmental sensors, etc.) as they move about. The user may actively direct these sensors, or allow them to operate automatically, providing a rich stream of information to be absorbed into the database. Furthermore, a user may develop higher level interpretations and understanding about the environment. The user can make precise observations and annotations based on what she sees. To help capture and share that information, the system must support annotation (the capability to annotate and edit the database) and collaboration

(the ability to communicate and work with remotely located users).

These high level capabilities and attributes translate into the following system design requirements:

Real-time Data Acquisition: To update the situational picture, the latest information must be placed in the database. Real-time acquisition and insertion of data must be possible from a variety of sources.

Dynamic Data: Because the database is continually being updated with spatial data, dynamic data structures must be capable of changing on-the-fly as pertinent information is received.

Synchronized Data Structures: Synchronized data structures are required, since there may be multiple, networked databases and multiple collection points.

User Geo-Location: A key situational variable is user location, so that users needing situational information can see who contributed it, and where it comes from. Users must be geo-located, and must be able to insert their orientations and movement vectors into the database. The database will be enriched with these annotations and, in turn, will permit more accurate user locations and orientations to be derived using sensor fusion techniques.

Global Geo-spatial Structure: There must be a global, geo-spatial structure capable of handling all types of geo-spatial data from all over the Earth. The database must be able to handle data from any location, and must therefore be robust enough to address any location at any reasonable resolution.

Time and Location-specific Services: This presumes a database structure appropriate to the location-specific and temporal nature of the user's situation and work. With this structure, one can have location-specific services, which also depend on the time at which relevant events occur, on what the user is doing, and on where she is and where she is going.

Situational Visualization can share many properties of Virtual Environments. The application may take the form of an immersive and highly interactive 3D virtual environment. Situational Visualization is also an Augmented Reality application in that one is simultaneously navigating the real world and a corresponding annotated 3D virtual replica of the world. There are issues of view registration in terms of aligning a virtual view in the database with the real view seen by the user. Although we do not do so in this paper, one could carry these issues all the way to the simultaneous alignment of virtual geo-located objects and real ones in the same scene, as is usually done in Augmented Reality. Because of the overlap, we treat the issues of immersion and interactivity in Situational Visualization from the standpoint of Virtual Reality.

This paper will examine the design of data and synchronization structures, as well as collaboration and awareness features to support Situational Visualization. We also examine some example Situational Visualization applications.

2. RELATED WORK

Some aspects of Situational Visualization are seen in systems that present information based on the user's location. A number of tour guide and map navigation applications have demonstrated location-specific information for small geographical regions. Georgia Tech's Cyberguide [1] is an indoor and outdoor tour guide based on the Apple Newton handheld computer. The indoor version presents information about various demos using an array of infrared transmitters for pinpointing location within a large lab. The outdoor version uses GPS to present information about local establishments in Atlanta. A travel diary is created and used to suggest other locations that a user might be interested in visiting.

The Touring Machine [7] is a wearable augmented reality system that presents multimedia information about the Columbia University campus with access to department web pages. It combines see-through augmented reality with tablet input and output. The "map-in-the-hat" system developed by Thomas, et al. [17] is a navigational aid system. It uses a see-through head mounted display to display waypoints and a compass for walking direction. The main focus of "map-in-the-hat" is to guide the users in an orienteering task from one point to another.

Nokia's Context Compass [16] is a wearable computer application that presents a see-through 2D overview map of an urban area with roads, buildings, and virtual objects. Virtual objects can represent buildings, other users, or messages, photographs, or waypoints, for example. Virtual objects are displayed on a linear compass and selected by facing the object, which centers the object in the compass. Each virtual object has an associated URL with a web page, image, audio, or video.

Our focus is to extend the notion of providing navigation and informational aids with real time data feeds and large data sets covering large (up to whole Earth) geographic regions. Various projects for spatialized hypertext such as WorldBoard [15] and the Real-World Wide Web [11] will utilize large information data sets over worldwide geographic regions, but our system provides an emphasis on situational awareness and visualization, requiring a more powerful set of geo-located objects and services than hypertext and other web media. It also permits continuous visual navigation to any point on the earth.

The collaborative and information annotation aspects of Situational Visualization are related to work at Columbia on indoor and outdoor interfaces for MARS [9]. This follow-on work to the Touring Machine adds a collaborative interface to their mobile augmented reality system that allows stationary remote experts to monitor and assist mobile users. The users can create virtual objects and annotations to point out interesting objects and direct attention. Another application that demonstrates context tagged annotations is StickePad on a GPS equipped PalmPilot, which allows users to create geo-located and time-stamped "stick-e notes" for ecology fieldwork [13].

Another project that is closely related to our description of Situational Visualization is the Battlefield Augmented Reality System (BARS) developed at the Naval Research Laboratory [10]. It also combines a detailed map of the user's surroundings overlaid with a context and situation aware information display. This AR system presents data about the soldier's environment such as sniper positions or

navigational waypoints based on the importance of each of the soldier’s current tasks and needs. However, it currently is designed to address the needs of the dismounted soldier in urban landscapes. We are developing a more general framework of services to assist a wider variety of users.

Our system is built on our global geo-spatial data organization and visualization system, VGIS (Virtual Geographic Information System) [3, 4, 6]. VGIS treats the terrain as a single, connected surface for rendering using a continuous, view-dependent Level-Of-Detail (LOD) representation. Although there are similar approaches in the literature, VGIS is the first system to support real-time view-dependent, pixel-accurate multi-resolution rendering of high resolution global datasets.

The VGIS system has the following capabilities:

- It accepts and integrates diverse types of geo-spatial data into a global framework.
- It is scalable to very large data stores and to distributed environments.
- It provides an interactive visualization framework so that the data store can be navigated as one coherent whole.
- It supports discovery of new information via navigation in context with unfolding detail.

Our approach starts with a hierarchical structure for optimal interactivity for data exploration. We use a "forest of quadtrees" with nested coordinate systems for handling global-scale terrain data and permitting timely paging of collections of objects in support of real-time navigation.

We have built a quadtree and shared cache data structure that together constitute the basic components of our global data model. The global structure is divided into 32 zones, each 45 degrees \times 45 degrees on the Earth’s surface. Each zone has its own quadtree and all quadtrees are linked so that objects or terrain crossing quadrant boundaries can be rendered efficiently. On top of this structure we have added 256 \times 256 coordinate systems within each quadtree – one coordinate system per node, eight levels below each root node – resulting in a 1 mm worst case precision.

3. DYNAMIC AND SYNCHRONIZED DATA STRUCTURES

An attribute of Situational Visualization is that data can be received all the time, either in discrete or streaming fashion. Our interactivity requirement means that these data must be readily available when received. One data structure that can effectively accommodate these attributes is an appropriately tailored universal hierarchy. In the case of geo-spatial data, the choice is the VGIS hierarchy described above, which we have shown can be tailored to handle a wide variety of data including terrain (from low to high resolution), phototextures and maps, GIS information, 3D buildings, and even time-dependent, volume data (e.g., weather) [3, 4, 6, 14]. The forest of quadtrees is used at the top level for each of these types of data and although each has its specific structure at higher levels of detail, the specific structures are quadnode aligned. The quadnode alignment permits the handling of dynamic data in the same general fashion, regardless of the type. As the data are received

they are organized to fit the quadtree hierarchy. This entails building links to data chunks at appropriate levels of the hierarchy and updating quadnodes with the appropriate averages, ranges, error values (for LOD management), and links. For streaming data this process may exceed the allotted time budget, in which case the dynamic data are organized approximately in the hierarchy. The system can then complete the organization during idle times.

For Situational Visualization, data must also be in "synchronized data structures" in the sense that any user may be collecting data and putting it in her own database, which must then be synchronized with databases of her peers. With geo-spatial data, one of the peers may be quite a bit larger than the others (e.g., a data server), so the networking structure and overall data organization will have aspects of both peer-to-peer and client-server structures. Our general, global structure allows us to think in terms of a complete global virtual hierarchy. This is the hierarchy for the whole Earth down to very high resolution (currently 1mm). No one, not even the servers, has a complete set of data; everyone has some piece of it. But if anyone receives a new piece of data for any place on Earth, he knows where to put it in the hierarchy and it is a unique location. The tree structure is unique but properties such as links to individual blocks of terrain are not. Thus one peer may have a set of elevation posts at a quadnode that differs from those at the same node of another peer. Furthermore, these node contents may change over time.

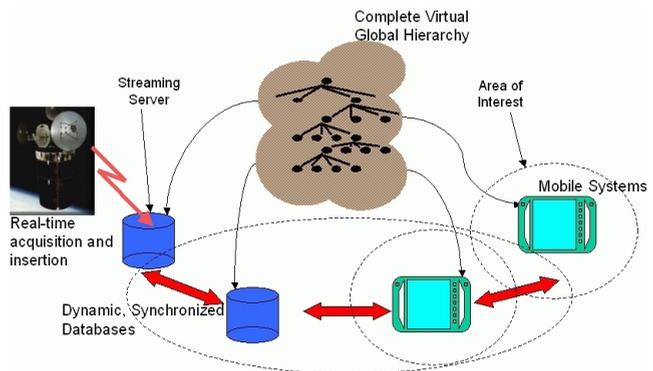


Figure 1: Distributed databases. The two mobile systems on the right have overlapping areas of interest. This figure is reproduced in color on page 201.

The distributed set of databases needs a synchronization mechanism. (see Figure 1.) If a mobile system, for example, moves into a new area and finds it doesn’t have the data it needs, it looks for these data amongst its peers (see Figure 1). On the other hand, it can also be collecting data that it wants to transmit to its neighbors. To support efficient synchronization we define 3 types of peers, each with different attributes: neighbors, servers, and collaborators. Neighbors are the set of all peers who may communicate; servers are neighbors with large repositories of data at defined locations, and collaborators are neighbors who are in active communication. Note that these groups are not mutually exclusive and could also change over time. The servers are the first systems queried for data in their areas of interest,

although the rest of the neighbors could also be queried for high priority requests. Collaborators directly share annotations or acquired data and also such things as interactions and viewpoints. Servers also get any updates for their areas of interest. As long as a peer gets a fresh download of data from the servers when it enters the collection of neighbors (but only for those regions in its area of interest that are marked newer on the appropriate server), it will be synchronized. Thus the synchronization is restricted to areas of interest for all peers. The meaning of areas of interest for mobile users will be discussed further below.

For efficiency the querying peer can construct an exact address for the data block it needs. This is one of the attributes of the virtual global hierarchy. For example, a terrain footprint can be addressed by first obtaining the zone it is in (i.e., the quadtree from the forest of quadtrees), then finding the leaf cell at the quadlevel appropriate to the footprint size. These are both bin sort operations and quite fast. A unique address can then be constructed. This querying procedure needs to take account of the distributed environment. Thus the footprint covers not only the scene in direct view but also areas likely to be covered in further navigation. This navigational footprint (and the resolution needed) will be different depending on, for example, whether a mobile user is walking, in a vehicle, or in an airplane. The system constructs two navigational footprints around the user's current location (see Figure 2). The outer one defines the region to be retrieved from the server, if necessary. However, this only happens if the user moves outside the inner footprint in her explorations. This arrangement insures that thrashing will be minimized as the user moves back and forth; once the user moves outside the inner footprint, a new set of footprints is set up centered around the new current position (Figure 2). This arrangement is similar to the one used in [12], except that now it is done in the context of a complete global hierarchy. The footprint address and resolution needed is then transmitted to, say, the relevant server. The server finds the appropriate parent node and then does a breadth first traversal across child nodes at the appropriate resolution. A list of data blocks is constructed for transmission (for example, terrain elevation blocks, or 3D blocks of buildings). The list is ordered so that blocks in the immediate vicinity of the user are sent first, minimizing the delay in receiving data of immediate interest. The global hierarchy permits these to be found quickly. The individual blocks are large enough that they contribute significantly to a given view yet small enough that they are transmitted quickly. Several blocks can be in a view, depending on the viewpoint location. We have found that this balance in size makes out-of-core data retrieval efficient for a variety of data [4, 3].

3.1 Streaming Data

Some of the data may stream in continuously. These may be time-stamped 3D weather data acquired from sensors such as Doppler radars, computer simulations, streaming video, or other data. A special server collects and organizes these data, but in addition, mobile users may also collect and transmit streaming data (e.g., video). The overall organizing structure for the streaming data is the same virtual global hierarchy discussed above. However, the data are also organized for fast queries of the form, "Tell me about this event at this time and location." To handle these queries we are in the process of developing an event-based hierarchy for

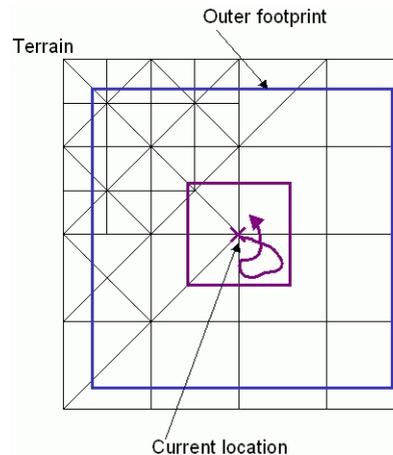


Figure 2: Inner and outer footprints used to manage retrieval of terrain and imagery data from a server. Spatial locality of reference helps minimize the amount of data that needs to be transmitted.

the sequence of steps in the streaming data (e.g., time steps or image frames) aligned with the quadtree structure. One issue we are studying is at what quad level (i.e., what level of spatial granularity) to insert the event structure. With this structure we will be able to ask questions such as, "How many tornadoes occurred in North Georgia during the Summer from 1995 to 1999, where, and with what severity," or, "At what times and in what locales did this atmospheric pollutant exceed this level during Spring and Summer, 1998." Queries such as this are relevant to the weather application involving streaming 3D Doppler radar data that we describe below. The above dynamic structure is discussed more fully in [8].

4. AWARENESS AND COLLABORATIVE STRUCTURES

Our implementation of Situational Visualization extends the spatial model of interaction proposed by Benford and Fahlen [2]. In the spatial model of interaction, there are three regions of note: the focus, aura, and nimbus. The focus is the area of interest, or region in which the user is aware of activity and objects. This could be the area defined by the limit of visibility. The aura is the region where interaction is possible. When the auras of two users intersect, the users can interact. For example, when two people are face to face and within a few feet of each other, they can easily converse. The nimbus is the region in which a second user can be aware of a first user. For example, on a crowded city sidewalk, a tall person has a much larger nimbus than a shorter person. They are easily noticed by others from longer distance.

There are four ways to modify the regions of focus, aura, and nimbus. The first is implicitly by making changes in location or orientation. The second is by explicit modification of key region parameters, such as radius. The third is by adapter objects, such as a telephone, telescope, or other real or imaginary object. The fourth is by boundaries in space such as walls.

This interaction model is employed in networked collaborative virtual environments such as MASSIVE and DIVE. However, there are several key differences between those virtual environments and a networked collaborative mobile environment, such as a Situational Visualization. The spatial model of interaction was developed for virtual environment where users have one virtual location and can easily move in virtual space. Some participants in a Situational Visualization are stationary clients. These clients could present an essentially unmodified spatial interaction model. However, many clients will be mobile. In these mobile clients, users will have both a real and a virtual location that may or may not coincide. The real or physical location corresponds to location reported by GPS. Virtual location corresponds to the user's cursor or viewpoint in the visualization. Since a user will often rapidly navigate through the visualization, virtual location will change faster than physical location. A user will need to be aware of events and objects surrounding their physical location, while they closely examine an area that may be a hundred miles away. Moreover, the user may wish to monitor several regions without explicitly examining those regions. A user may also want to be notified of future events that may have significance.

Modifications to the concept of focus would be to create focus (or awareness) regions for both the user's real and virtual locations. The user would also be able to define additional regions of focus, in a sense, dropping pieces of awareness across the terrain or attaching them to particular users or moving objects. Users who make modifications to the database should transmit updates to other users who have an overlapping focus region. Since the spatial model of interaction defines different region extents for different media (sight, sound, communication, etc.), it follows that users should be able to register focus for particular classes of objects as well, (weather, traffic, virtual objects, other users, etc.) To deal with warnings about future events, a prediction process would determine if objects, such as weather events might intersect a focus region in the future, even if the focus region is in motion with the user.

In our implementation, auras determine interaction such as personal communication and database communication. Since physical location typically changes less rapidly than virtual location, the concept of aura may be split in two. Transmitting auras will be attached to virtual body and receiving auras will be attached to physical body. When a user wishes to initiate contact, the user will navigate the virtual body in proximity of another user's physical body and establish a two-way communication channel. Since users may wish to stay in contact even if users move apart, the channel must be explicitly closed. In our implementation, personal communication is currently text messaging, with voice communication to be implemented in the near future. Auras will also apply to database communication. Some users may carry updated static terrain, building, or other virtual object information. For example, one user may have an annotated database with locations of wheelchair ramps, handicapped parking spaces, automatic doors, and stairways. A second user can initiate a database synchronization in much the same way as a communication channel. This allows the second user to fetch and add all the new information from the first user's database. The second user can also have a previously enabled database connection that fetches this information from the first user as needed. Some special pur-

pose stationary servers such as a weather information server would project a location quite different than its physical location. Its database communication aura coincides with the weather footprint.

Mobile Situational Visualizations will have nimbus regions for virtual location as well as physical location. Nimbus is the visibility profile of a user. In a standard mode of operation, users would only be aware of the nimbus surrounding the physical location of other users. Users would only see the physical locations of other users. If a group of users were collaborating closely, they may wish to see an avatar representing a user's virtual location.

5. SITUATIONAL VISUALIZATION APPLICATIONS

There are many potential Situational Visualization applications. One class has to do with awareness, being notified of time-evolving events based on one's location, direction of movement, or actions. The nano-forecaster described below is one such application. Another would be a "traffic-wizard" that gives one advice about traffic on potential routes as one approaches. Another class would involve geo-referenced services, such as telling a mobile user where and in what direction restaurants, bus stops, theaters, or other services are based on her current location. Other applications could involve annotations and updates of the synchronized database as determined by real world observations. These might include finding and noting handicap access points, noting positions of trees or shrubs and then placing appropriately scaled models, noting incompleteness or inaccuracies in building models and then capturing images to be linked to the models, capturing geo-located and oriented video, and many other applications. There are also collaborative applications, such as two mobile users investigating a large building site and coordinating their information collection so that it is complete. Other collaborations might involve mobile users and those at some central location or command post. The user at the central location might act as an overall coordinator, directing the mobile users and sending or receiving relevant information.

Our prototype Situational Visualization system consists of a set of central computers and a mobile system. The central computers include both desktop machines and a virtual workbench. The latter enhances collaborations since multiple people can be at the workbench and more than one can interact. All systems are connected by an 802.11b WaveLAN wireless network. To reach mobile users, WaveLAN antennas have been mounted on the exteriors of two buildings (Figure 3) providing a communication path of about a half-mile through a central part of the Georgia Tech campus. A nominal 11 Mb Ethernet runs on the wireless network. The outdoor portion of the network is subject to interruptions and points where communication may be blocked. This, however, permits us to try out predictive tracking, use of cached events, or other ideas for handling interrupted service. The mobile computer is an IBM 850 MHz Thinkpad with 512 MB RAM and an ATI Mobility 128 AGP graphics chipset with 16 MB video RAM. This configuration runs the VGIS global geo-spatial visualization system at interactive rates.

To evaluate the Situational Visualization concept and begin building the needed infrastructure, we have identified

Georgia Tech wireless testbed



Antennae (approx. 0.5 miles apart)

Figure 3: The Georgia Tech wireless testbed for Situational Awareness. To reach mobile users, WaveLAN antennas have been mounted on the exteriors of two buildings providing a communication path of about a half-mile through a central part of the Georgia Tech campus. This figure is reproduced in color on page 201.

two prototype applications. The first is an awareness application, the Nano-Forecaster. The second is an annotation application, the Wheelchair Assistant.

5.1 The Nano-Forecaster

Up to now the best weather reports that a person can receive is a set of "micro-forecasts" developed for various areas of a geographic region. These are not specifically correlated to a viewer location and do not provide accurate information such as "When is a storm going to hit me, from what direction, and with what intensity?" However, such information can be made available to a user if integrated in the above Situational Visualization context with appropriate geo-spatial model and weather data. We have an accurate model for North Georgia with high resolution insets for parts of Atlanta including the Georgia Tech campus and the weather data (in this case, real-time streaming 3D Doppler data from multiple radars covering the northern half of Georgia). Similar models have been or could be developed for other locales. We have previously shown how 3D, time-dependent data can be inserted on-the-fly into the geo-spatial hierarchy [8]. The Doppler radars are operated by the National Weather Service and are distributed across the country, especially around urban or tornado-prone areas. Thus the components are available that could make possible wide use of the nano-forecaster.

The mobile visualization system is shown in Figure 4, where the user has navigated to a ground level view associated with her location. We have found that a GPS-controlled icon, as shown in Figure 5, gives reasonably accurate and easy to follow position updates in an overview as a user walks through the same area. Further, the system can be made to automatically navigate to a view encompassing the user's current position if its viewpoint is initially somewhere else. The moving position information is enough to dynamically update nano-forecaster queries. Since the time-dependent, 3D Doppler data are in the global geo-spatial

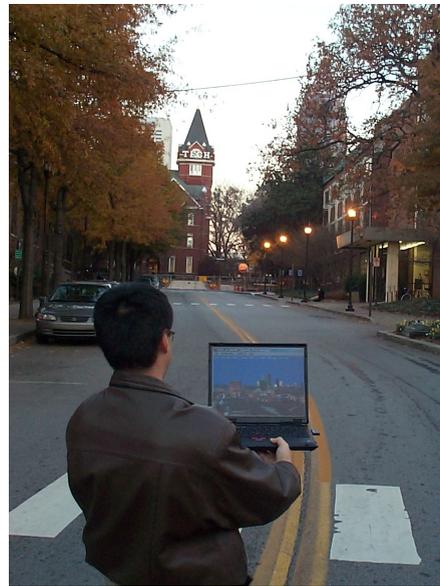


Figure 4: Prototype laptop Situational Visualization system. Mobile VGIS is running on the laptop, showing a location on Georgia Tech campus plus buildings in midtown Atlanta. This figure is reproduced in color on page 201.

hierarchy, the parts in the user's current area of focus can be quickly located on the streaming data server. Associated with these 3D data is a set of storm analyses that are generated at the same time as the data are collected. These track windfield intensities, reflectivity, and spectral width. These quantities can be transformed into severe storm signatures (including tornadoes), and precipitation intensity and type (e.g., rain or hail). For signature above certain values, features called mesocyclone cells are generated, which have a location, width, height, and average values for the signatures. These cells move over time and can be tracked. Certain cells that have the signatures of tornadoes, for example, can be flagged. The tracking permits a simple, but accurate phenomenological model to be applied for predicting cell motion. A severe storm passing over Atlanta, with both 3D radar data and mesocyclones, is shown in Figure 6.

In the current nano-forecaster prototype, the forecasting agent looks only at the subset of mesocyclones, with their predicted paths, in the current area of focus. It determines whether the path of any mesocyclone will pass over the user's projected position (taking into account trajectory information from previous positions). The area of focus includes a time frame of interest; if the intersection is within that time frame and the mesocyclone signatures are above a selected threshold, the user is notified. Notification can be as simple as an arrow, indicating the direction of the approaching weather, labeled with type and severity of storm and time of arrival, as shown in Figure 7. Alternatively, the user could request a detailed view of the 3D Doppler data in her area of focus, as in Figure 6. The user's area of focus and trajectory and the storm's development are all continuously updated. Since the key notification involves very little data, it could even be displayed on a PDA with minimal display capabil-

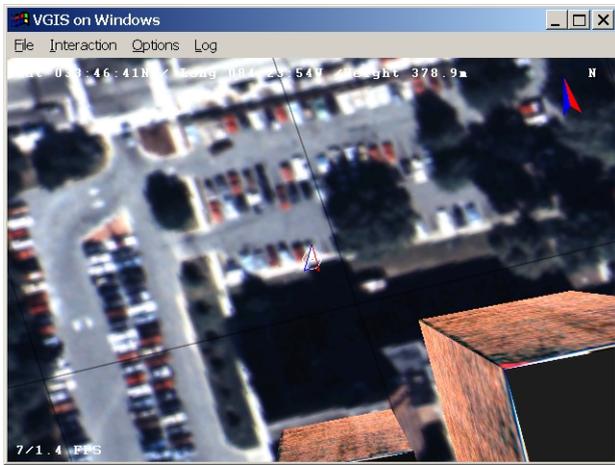


Figure 5: Screenshot of Mobile VGIS tracked by GPS. The arrow at the center indicates the user's real position. This figure is reproduced in color on page 201.

ities, perhaps just showing a fixed image of the user's area with her current position and the storm's labeled vector superimposed.

5.2 Wheelchair Assistant

The Wheelchair Assistant is an application that allows wheelchair users to navigate appropriately through an outdoor urban landscape. Users can annotate the building and terrain database with icons denoting locations of wheelchair ramps, handicapped parking spaces, automatic doors, and stairways. Users will avoid having to double back because of an outdoor staircase or use a heavy door when an automatic door is readily nearby.

The user can add wheelchair annotations by selecting an annotation type and clicking in the VGIS window for the location. The user can also move to the proper location and then select an annotation type. After a short delay, the annotation will be made. One of the advantages of displaying a highly detailed set of satellite phototextures and 3D building models, as in VGIS, is that a user can accurately place an annotation even though GPS location information is often inaccurate. The user can input the position of a wheelchair ramp in relation to the building models or items depicted in the satellite photo. Furthermore, the user can adjust the GPS position to more accurately reflect the user's actual position.

Through database synchronization, users can share these annotations to create a complete picture of wheelchair access for an area. This allows other wheelchair users to benefit. It also allows an organization like a business, school, or city to visualize and assess how wheelchair users are constrained or assisted.

A similar application mapping locations of emergency telephones, police stations, and manned security desks might also be a part of crime prevention and assessment projects.

6. CONCLUSIONS

In this paper, we have introduced a new style of visualization called *Situational Visualization*, based on our mo-

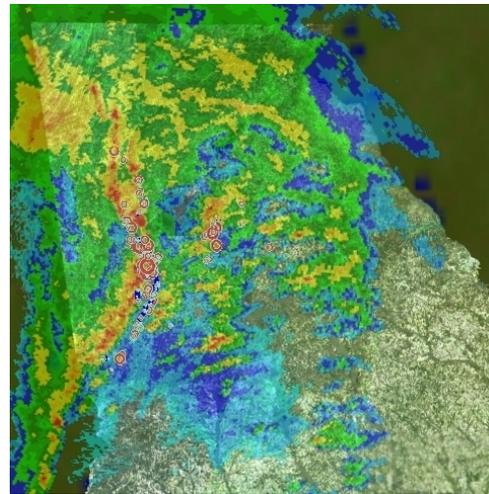


Figure 6: Overview of Doppler Radar readings over North Georgia. The column of cylinders near Atlanta to the upper left is a series of mesocyclones, or large circular wind patterns. This figure is reproduced in color on page 201.

bile Virtual Geographic Information System, VGIS. Our approach to issues surrounding how to deploy augmented information spaces to mobile users is to start from the basis of a rich source of information about the real world, such as terrain elevations and aerial photography supplied by various government agencies. This allows us to build an augmented information source that already has a great deal of ground truth. This Situational Visualization system allows the user to enhance the visualization and any underlying simulation by inputting the user's observations of the phenomena of interest, thus improving the quality of visualization for the user and for any other users that may be connected to the same database. The example applications show how such a system can afford a rich information space that can be used to enhance the user's situational awareness.

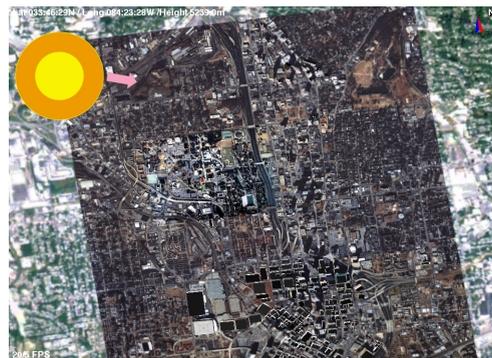


Figure 7: Closer orbital view of Georgia Tech and midtown and downtown Atlanta with a single mesocyclone plus its predicted path over the next 5 minutes. This figure is reproduced in color on page 201.

7. ACKNOWLEDGMENTS

This work is supported by a grant from the NSF Human-Computer Interaction Program and the DoD MURI program administered through the Army Research Office. This work is also supported by a grant from the NSF Large Scientific and Software Data Visualization Program.

8. REFERENCES

- [1] G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton. Cyberguide: A Mobile Context-Aware Tourguide. *ACM Wireless Networks*, 3(5):421–433, 1997.
- [2] S.D. Benford and L.E. Fahlen. A Spatial Model of Interaction in Large Virtual Environments. In *Proceedings of Third European Conference on CSCW (ECSCW'93)*, pages 109–124, September 1993.
- [3] D. Davis, T.Y. Jiang, W. Ribarsky, and N. Faust. Intent, Perception, and Out-of-Core Visualization Applied to Terrain. *IEEE Visualization '98*, pages 455–458, October 1998.
- [4] D. Davis, W. Ribarsky, T.Y. Jiang, N. Faust, and Sean Ho. Real-Time Visualization of Scalably Large Collections of Heterogeneous Objects. *IEEE Visualization '99*, pages 437–440, October 1999.
- [5] M.R. Endsley. Situation Awareness Global Assessment Technique (SAGAT). In *Proceedings of the IEEE 1998 National Aerospace and Electronics Conference*, volume 3, pages 789–795, May 1998.
- [6] N. Faust, W. Ribarsky, T.Y. Jiang, and T. Wasilewski. Real-Time Global Data Model for the Digital Earth. In *Proceedings of the International Conference on Discrete Global Grids*, March 2000.
- [7] S. Feiner, B. MacIntyre, T. Hollerer, and A. Webster. A Touring Machine: Prototyping 3D Mobile and Augmented Reality Systems for Exploring the Urban Environment. In *Proceedings of the First International Symposium on Wearable Computers (ISWC 1997)*, pages 74–81, October 1997.
- [8] B. Hannigan, R. Mitchell Parry, Nickolas L. Faust, T. Y. Jiang, and William Ribarsky. Hierarchical Storage and Rendering of Real-Time 3D Data. In Nick Faust, editor, *Visualization of Temporal and Spatial Data for Civilian and Defense Applications III*, volume 4368A, 2001.
- [9] Tobias Hollerer, Steven Feiner, Tachio Terauchi, Gus Rashid, and Drexel Hallaway. Exploring MARS: Developing Indoor and Outdoor User Interfaces to a Mobile Augmented Reality. *Computers and Graphics*, 23(6):779–785, 1999.
- [10] Simon Julier, Marco Lanzagorta, Yohan Baillet, Lawrence Rosenblum, Steven Feiner, Tobias Hollerer, and Sabrina Sestito. Information Filtering for Mobile Augmented Reality. In *Proceedings of IEEE International Symposium on Augmented Reality (ISAR 2000)*, pages 3–11, October 2000.
- [11] Rob Kooper and Blair MacIntyre. The Real-World Wide Web Browser: An Interface for a Continuously Available, General Purpose, Spatialized Information Space. In *Proceedings of the 2nd International Symposium on Mixed Reality*, March 2001.
- [12] Michael R. Macedonia, Michael J. Zyda, David R. Pratt, Donald P. Brutzman, and Paul T. Barham. Exploiting reality with multicast groups. *IEEE Computer Graphics and Applications*, 15(5):38–45, 1995.
- [13] J. Pascoe. Adding Generic Contextual Capabilities to Wearable Computers. In *Proceedings of the Second International Symposium on Wearable Computers (ISWC 1998)*, pages 92–99, October 1998.
- [14] Christopher D. Shaw, R. Mitchell Parry, William Ribarsky, Anthony A. Wasilewski, and Nickolas L. Faust. Interactive Volume Rendering of Non-Uniform Data. In Nick Faust, editor, *Visualization of Temporal and Spatial Data for Civilian and Defense Applications III*, volume 4368A, 2001.
- [15] J. C. Spohrer. Information In Places. *IBM Systems Journal*, 38(4):602–628, 1999.
- [16] Riku Suomela and Juha Lehtikoinen. Context Compass. In *Proceedings of the Fourth International Symposium on Wearable Computers (ISWC 2000)*, pages 147–154, October 2000.
- [17] B. Thomas, V. Demczuk, W. Piekarski, D. Hepworth, and B. Gunther. A Wearable Computer System with Augmented Reality to Support Terrestrial Navigation. In *Proceedings of the Second International Symposium on Wearable Computers (ISWC 1998)*, pages 168–171, October 1998.