Visualizing Real-Time and Archived Traffic Incident Data

Michael VanDaniker
University of Maryland, College Park
Center for Advanced Transportation Technology
College Park, MD 20740
Phone: (301) 403-4628
Fax: (301) 403-4591
Email: mvandani@umd.edu

December 23, 2008

ABSTRACT

The sheer volume of data associated with a single transportation incident is difficult to fully comprehend, and existing incident management systems often fail to equip users with the necessary tools to make sense of this data. The consequences of delayed decision making, misinterpreting data, or even overlooking pertinent information can be life threatening. A tool that gives the incident manager the ability to grasp the entire picture of an incident in a minimal amount of time is essential to facilitate quick and accurate decision making.

The Transportation Incident Management Explorer (TIME) provides for the visualization of real-time and historic traffic management center incident data. All data pertinent to an incident can be viewed in a compact graphical overview. Temporal data are rendered as timelines where color and length convey information that would otherwise take significantly more space to display and more time to understand. The mouse can be used to scroll over elements in the timeline to glean additional information. Geospatial data are plotted on an interactive map, providing users freedom to explore regions affected by an incident. TIME reduces the chance of missing critical information, enables users to correlate events, and quickens the time needed to comprehend the many simultaneous events that occur or occurred during the management of an incident. This paper describes the use of the TIME tool when visualizing real-time and archived incident data collected and provided by the Maryland CHART, Washington, D.C., and Northern Virginia Traffic Management Centers.

INTRODUCTION

Traffic management centers around the world use real-time incident management systems to record and coordinate emergency response to transportation events. The records produced by these systems include the time and location of the incident, the number of vehicles involved, agency notifications, lane closures, detector measurements, and many other properties

and their effects on traffic conditions. Proper decision making depends on the traffic incident manager's ability to rapidly understand the complete story of an incident, but existing tools have failed to make this an easy task. Most incident management systems require an operator to navigate through several pages of free text, tables, graphs, and maps to even begin to get an idea of what is being done to manage an incident. This process can be time consuming and error prone. A tool that mitigates these two problems can streamline the incident management process and lead to faster accident scene clearance times and ultimately to safe roads.

This paper presents the Transportation Incident Management Explorer (TIME), a visualization that combines the temporal and spatial data associated with incident logs, presenting, on a single webpage, all the events recorded during incident management. TIME is meant to help incident managers quickly assess a situation and make informed decisions. The system works with live data, providing up-to-the second visualizations of an active incident, but it can also load archived data, facilitating after-action reviews and allowing individuals to quickly understand the consequences of the actions taken in response to a past transportation incident.

The Nature of Incident Management Data

Incident management records contain several *types* of data. These include:

1. Temporal data
    a. Discrete time points: Events that occurred at an instant in time.
    b. Time-intervals: Events that have start and end times.
2. Spatial data
    a. Geographic data: Any items that contains latitude and longitude coordinates.
    b. Relative positioning data: While relative positioning can be derived from latitude and longitude, certain items in the dataset only include information about their positioning relative to similar items.
3. 1-Dimensional data: Data that are described by plain text or numerical values and have no temporal or spatial attributes.

These data *types* manifest themselves throughout the dataset in the various *categories* of data recorded by incident management systems. These *categories* include:

1. **Communication Logs:** The communications between individuals managing an incident scene occur at discrete points in time.

2. **Variable Message Signs:** Incident managers can post messages on variable message signs (VMS) located alongside highways. These messages can be active for arbitrary time intervals, and the signs are geolocated.

3. **Responders**: Agency response records consist of both discrete and interval time data; Incident managers note the time an agency is notified and the window of time when an agency's dispatched unit is on scene.

4. **Lane Status**: Records of lane openings and closures are represented as time intervals. The configuration of the highway (number of on-ramps, shoulders, etc.) also provides relative spatial information about the lanes.

5. **Detector Data**: Detectors along the road record the number of vehicles passing by and their average speeds. Detector readings are coupled with the their absolute location on the road and their relative position with other detectors based on mile-markers.

6. **Miscellanea**: Various 1-dimensional data is also stored by incident management software. Among these records are the number of vehicles involved in the incident, the weather conditions when the incident occurred, the severity rating assigned to the incident, and a number of other fields.

Existing Systems

The end goal of the most work in this domain is to foster a safer transportation system. Many groups have attacked this problem from different directions. Road assessment programs *(1)* attempt to analyze problem spots along a roadway and re-engineer them to attempt to reduce the likelihood of future accidents. First responder training programs *(2)* educate potential

responders about how to place their vehicles, set up traffic control, and quickly clear a scene in order to minimize the impact of an accident on traffic flow. These training courses can involve simulating an accident scene based on data collected from other accidents. Because it is often prohibitively expensive to gather a number of first responders in one location and set up a safe yet realistic accident scene, there has been some attempt to create accident scenes in a virtual world that responders can log into remotely *(3)*. All of these systems incorporate data logged by incident managers on some level.

Despite the fact that incident management data is essential to all of these other programs, there are currently no software packages that can claim to be incident management *visualizations*. While many incident management *systems* exist, for the most part these tools are little more than a series of electronic forms which support data entry and device management. The emphasis is on data input, with less consideration given to full situational awareness and review capabilities. Regardless, by reviewing the contents of these forms, incident managers can learn what is currently being done to manage the incident scene. Unfortunately, the structure of these forms often makes assessing current conditions an arduous task. Doing after-action review from within the live system is often impossible because the forms only reflect the current conditions. Many systems try to support after-action review by allowing users to print out text-based history logs that list every action that was taken while managing an incident. These logs are often several pages long, and extracting the interesting and meaningful events is tedious.

Before development on TIME began, several incident management systems from the states of Maryland, Virginia, Georgia, and Texas were reviewed and their strengths and weaknesses were evaluated. The majority of the design decisions made during the development of TIME were influenced in some way by these evaluations.

All of the incident management systems surveyed are comparable in how they allow users to enter and, by extension, view data. Users enter 1-dimensional information (number of vehicles involved, weather conditions, etc.) using standard form controls like drop-downs and

check boxes, and they specify geospatial information, through some means, via a map. These controls in and of themselves provide all of the relevant information about their respective fields: Users can see what the weather conditions were at the time of the incident by simply reading the text in the appropriate drop down.

Unfortunately, these systems do not handle temporal data as elegantly as they handle spatial and 1-dimensional data types. Responders, for example, have a notification time, an arrival time, and a departure time. Many systems opt for a tabular layout when displaying this information which makes deciphering when the associated events occurred laborious. During after-action evaluation an incident manager might need to determine which responder was on-scene the longest. In a few of the incident managements systems reviewed, the duration is not made visible, and users would have to tediously subtract the arrival and departure times for each responder. Tables that display variable message sign information suffer from even more problems. These tables often have a column displaying the current message on a VMS, but they do not provide a way for users to view previously posted messages, even though this may prove helpful when determining the next update to post on a sign.

With respect to layout, most existing incident management systems suffer from one of two distinct classes of problems. Some systems attempt to fit all of the form elements relevant to an incident on a single screen which makes the user interface cumbersome and cluttered. Users are forced to scan the screen which is usually a sea of text and form controls for relevant information. As seen in Figure 1, other systems segregate their controls across several pages. In this multi-page layout the individual pages can be relatively sparse, making it easy for users to find information. Unfortunately users have to navigate between pages when trying to assess information that is distributed across the application. This takes up valuable time, and since users will likely only bring up the page that they think is relevant, important information hidden on the other screens may be overlooked. A system that displays all information on a single webpage and provides a scroll bar to reach controls that do not fit elegantly onscreen combines

the best features of these two layout techniques. This design allows the developer to carefully configure the positions of controls in each section of the system to avoid clutter while eliminating the need to navigate between different pages when assessing information. Although some information may be off-screen at any given time, it takes less effort to retrieve this information in a scrolling application than in a multi-page one because no data needs to load from the server. As users scroll between two sections, they will be exposed to the information in all of the sections placed between their initial position and their destination. A multi-page application does not allow for this sort of passive exposure; users must actively hunt for data. A cluttered single-screen application can cause information overload, preventing users from taking in specific pieces of information.
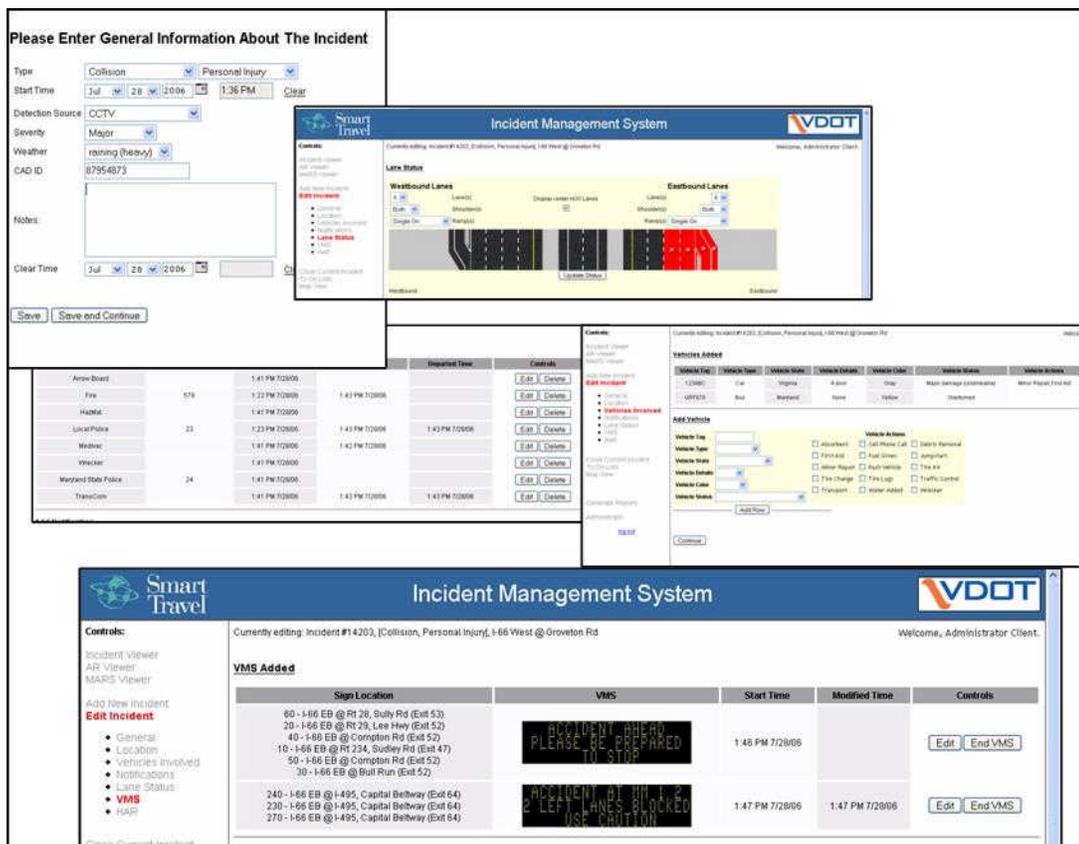

Figure 1: Five of the incident entry pages found in VDOT's Incident Management System.

RELATED WORK

The prominent features of a dataset often suggest appropriate visualization techniques. Nearly every piece of information in the incident logs has some sort of temporal data associated with it, so an interface that emphasizes temporal information may be beneficial. Timeline visualizations have been used for centuries to display the progression of time and are a good starting point when considering potential visualizations for temporal data. Although numerous temporal visualizations exist, many follow a common practice: plotting time-oriented data as horizontal lines across a temporal axis. Techniques that use the approach are hence forth referred to as "traditional timeline visualizations."

Careful consideration must be given when determining how to differentiate the various uses of time while still maintaining a consistent presentation. Aigner et al. catalogued several techniques used to visualize data based on how temporal information is used within a dataset *(4)*. Among the topics discussed are techniques for visualizing time intervals, discrete events, and cyclic and linear time. The traditional timeline visualization influenced the majority of the applications Aigner et al. surveyed.

Timeline visualizations have been used to successfully convey time-series data in a number of domains. PlanningLines *(5)* displays the duration of tasks in a project and includes methods for visualizing uncertain timeframes. These techniques are useful when differentiating events that have recently ended from events that are still on-going. ThemeRiver *(6)* displays temporal data differently; by plotting time along a horizontal axis and numerical values along a vertical axis, it allows for the search of trends in multiple time-series. Another application, TimeSearcher *(7)*, provides an environment for exploring generic time-series data. Timeline *(8)* is a project from the SIMILE suite of web-based visualizations that presents a customizable timeline widget that can be embedded in any webpage. The system is designed to display events occurring over an arbitrarily long time period. Users can drag the interface right or left to explore data occurring in the past or future, respectively. With LifeLines2 *(9)* users can align events based

on when they occurred relative to one another. Because most other timeline visualizations align events relative to a fixed timeframe, this method provides a novel perspective on the data and can lead to insights as to how events relate to each other.

The history flow visualization, developed by Viégas, Wattenburg, and Dave *(10)*, is used to display modifications to a body of text by multiple individuals. Revisions are rendered as vertical strips with color patches where each color corresponds to the individual who authored that portion of the document. When sequential revisions are aligned side by side, trends in authorship become apparent. The framework could be adopted to show edits to an incident's log by multiple individuals, and thus provide a method for holding users accountable for their actions. Although this is a compelling idea, it is left for future work.

Because of their long history and ease of use, traditional timeline visualizations are commonplace outside of academia. The media has long used timelines to convey complex temporal datasets to the public *(11, 12)*, and commercial calendar applications *(13, 14, 15)* use variations of the classic timeline to show the duration of scheduled events.

Luz and Masoodian *(16)* observed that traditional timeline visualizations that allocate a horizontal strip for each element in the dataset can rapidly become unwieldy because of their increasing height and wasted space. Their temporal mosaic (Figure 2) is an alternative to the traditional timeline that assigns vertical space such that each element is sized proportionally to the total number of active elements at a given time. The temporal mosaic works well with interval data, but it is not clear how it can be adopted to handle instantaneous events. It uses a space-filling algorithm, so a method for allotting a reasonable amount of space for instantaneous events needs to be devised before the mosaic can adequately handle the discrete time points found within the traffic incident dataset.
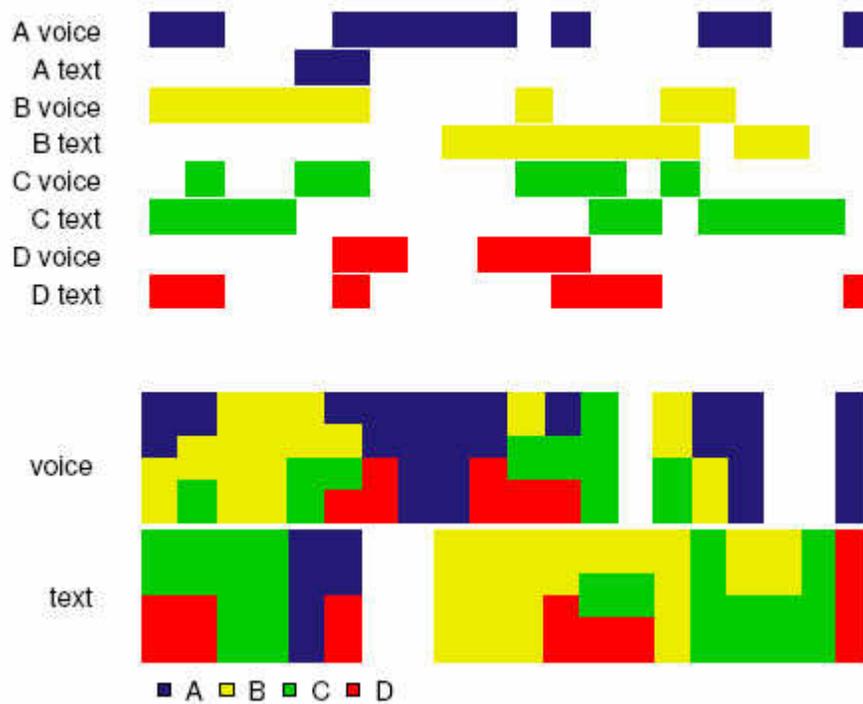
Figure 2: A traditional timeline and a temporal mosaic for the same dataset as they appear in Luz and Masoodian's paper on the temporal mosaic *(16)*.

The original LifeLines project *(17)* displays an individual's medical history using a series of stacked timelines to group related events, such as hospitalizations or the results of blood tests. Each item in the stack can be collapsed to free up screen real-estate, effectively combating the increasing height that Luz and Masoodian's temporal mosaic addresses. In traditional timeline fashion, each event is plotted as a horizontal bar with length proportional to that event's duration. Medical records can span several years, so controls are provided to allow users to zoom in on a time range of interest. The datasets visualized by LifeLines and TIME are both rich with various forms of temporal information. Because of this similarity, TIME borrows heavily from the lessons learned during the LifeLines project.

Although LifeLines uses medical records as an example dataset, the tool is meant as a "general visualization environment for personal histories" *(17)*. As a generic temporal visualization, it deals with data that is, for the most part, purely time-oriented. In addition to

time-oriented data, the traffic incident dataset contains a significant number of data elements that are qualified with spatial information (detector locations, VMS locations, etc.). Despite the effectiveness of existing timeline visualizations, many do not account for this interplay between temporal and spatial data.

While focusing primarily on spatial data, commercial GIS software packages *(18, 19)* address the need to integrate temporal data as well. GeoVISTA *(20)* uses animations to help users make sense of how data changes over time. Unfortunately, the timelines are rendered separately from the map, giving the spatial and temporal properties of the dataset a disconnected representation.

GeoTime *(21)* integrates spatial and temporal data with a 3D interface that plots position on the xy plane and uses the z axis to represent the passage of time. By tracing the path of objects in the positive z direction, a user can follow those objects as they move forward in time. This works well for visualizing the changes in position over time but needs to be expanded upon in order to account for changes in other properties, like the messages placed on the VMS.

Andrienko and Andrienko have long been proponents of visualizations that combine spatial and temporal data. In *(22)* they present a system that partitions a map into cells, preforms aggregation on the data within each cell, and renders two-dimensional plots for that data. Because of small geographic changes on the local level, one would often expect neighboring cells to have similar looking plots. When this expectation is not met, the user might find an interesting anomaly in the data in the affected regions. In *(23)* the authors present a number of techniques for overlaying time series data on a map when data is much less evenly distributed. These techniques manifest themselves in TIME in the form of timelines displayed on the map representing changes to variable message signs.

DESCRIPTION OF THE INTERFACE

The TIME visualization tool, shown in Figure 3, derives several of its visualization techniques from the traditional timeline. To separate different temporally-oriented data, TIME employs the stacking technique used by LifeLines. The TIME interface includes six visualizations:

1. Communications

2. Variable Message Signs

3. Responders

4. Lane Status

5. Traffic Speed

6. Traffic Volume

TIME adds the ability to drag and drop any of these six categorical visualizations (also referred to as "regions" of TIME) to rearrange them as needed. Each region may also be collapsed by clicking on its header. The collapsing process is animated to prevent users from becoming disoriented by rapid changes in the layout of the interface. The 1-dimensional data associated with the incident can be viewed as a tooltip by rolling the mouse over the blue information icon.
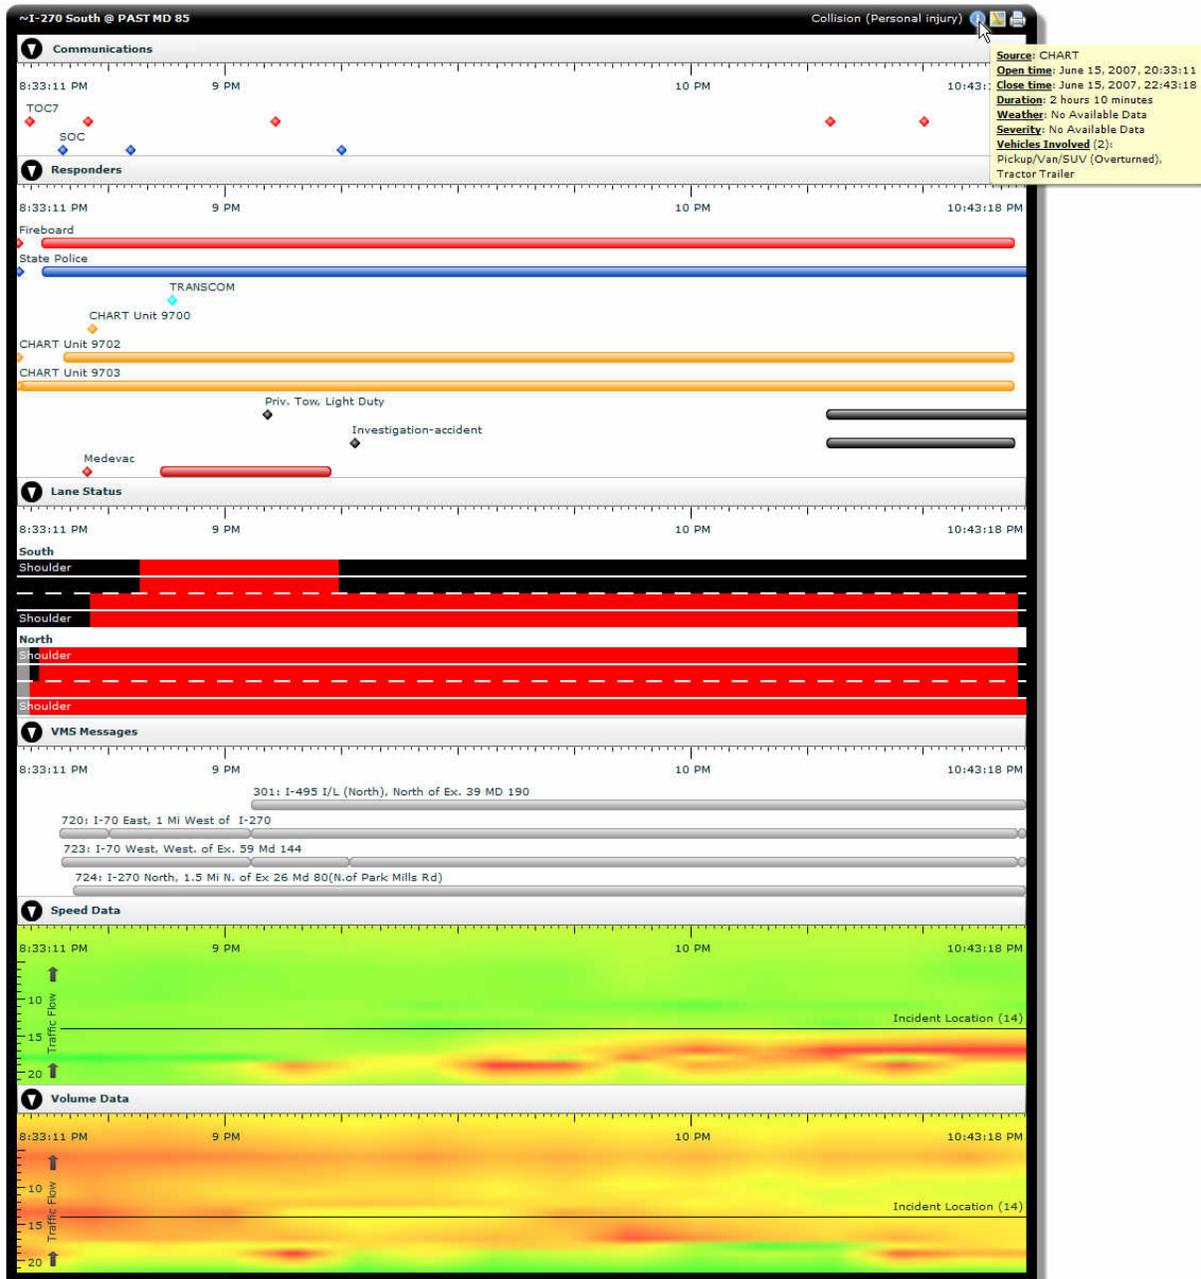
Figure 3: The Transportation Incident Management Explorer

The Interactive Timescale Control

In each category of data, temporal information underpins the visualization of choice. This is an important property of the dataset, and to make this information apparent to the users,

identical interactive timescale controls are placed at the top of each visualization (Figure 4).
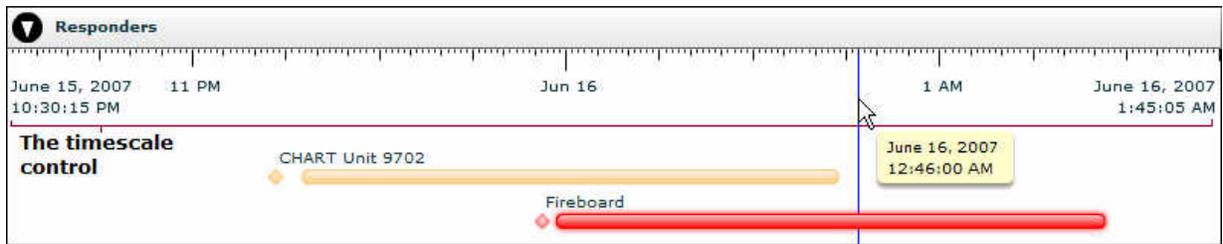

Figure 4: The timescale control

The layout of the timescale has been designed to convey as much information as possible without cluttering the interface. The start time of the incident is printed on the left of the timescale and the current time is printed on the right. In the event that the incident being viewed has already ended, the "close time" is used in place of the current time. A series of tick marks runs between the two ends at each minute mark. In the case of long-running incidents, placing tick marks every minute might result in multiple tick marks being drawn at the same position. If this were to occur, a five minute interval is used instead. This algorithm automatically scales up from five minute long, fifteen minute long, hour long, and day long intervals until an appropriate tick mark interval is found.

Like the markings on a yardstick, each tick mark is sized depending on its significance; one minute marks are the shortest, five minute marks are slightly longer, etc. Labels are placed at every hour tick mark, space permitting. By providing labels and a visual differentiation between the significance of each tick mark, users should be able to find arbitrary times on the timeline with less effort than they would when using a less descriptive timescale.

Tooltips indicating the time represented by each tick mark can be shown by rolling the mouse over any point on the timescale. In addition, a vertical line is drawn across all regions of TIME at the mouse's x position. Events that intersect the vertical line are emphasized while events that do not are grayed out. This draws attention to events occurring at the time under the

mouse cursor, making it easy for the user to identify everything that occurred at any given time.

By dragging the left and right borders of the tool users can expand the timeline, giving more screen real-estate to each time interval. This has the advantage of keeping the entire history of an incident visible while making more fine-grained details available. Because traffic incidents generally last no more than a few hours, a standard 1024 pixel wide screen can usually accommodate a sufficient level of detail for each event. Expanding a timeline for an incident lasting 120 minutes to 1024 pixels in width will allocate 8 horizontal pixels for each minute.

Data-Specific Visualization Techniques

The six categorical visualizations in the dataset have different properties associated with them, and TIME uses separate visualizations to accommodate these differences. Each technique draws on the traditional timeline. Instantaneous events are rendered as diamond-shaped markers while time-intervals are rendered as horizontal bars. Several existing temporal visualizations use bars for both instantaneous events and time interval events, but the diamond shape makes it easier to differentiate instantaneous events from relatively brief time intervals.

Communications

Incident managers who send a messages during the course of an incident are assigned their own horizontal strip spanning the width of the temporal axis (Figure 5). Because communications occur at discrete times, diamond-shaped markers are placed along the line, corresponding to the time when communications were sent. The tooltips for each marker display the time and text of the relevant message. TIME's communications visualization shows all messages regarding an incident, allowing users who were initially "out of the loop" to review the messages they missed.
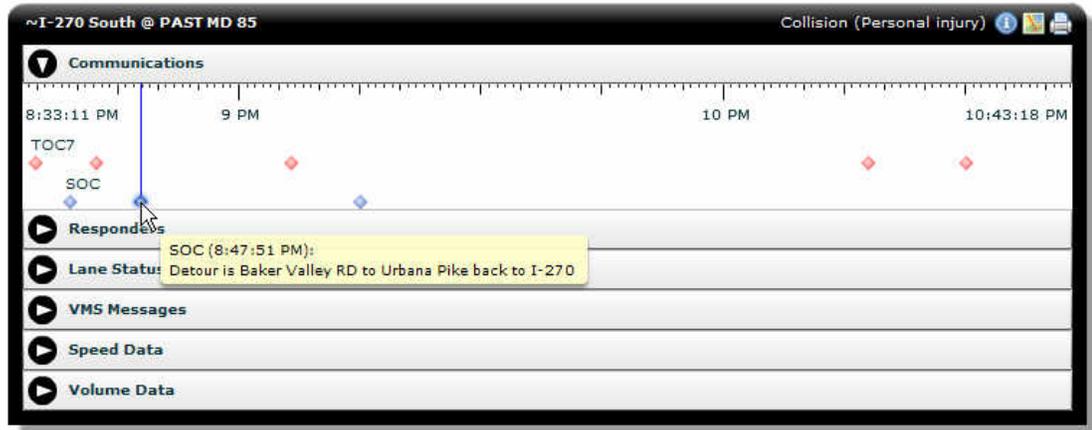
Figure 5: The communications visualization

Variable Message Signs

The timeline-based VMS visualization used by TIME allows users to see the full history of the VMS message postings, and its layout allows users to quickly ascertain the meaning behind the temporal data (Figure 6). Each VMS is rendered on its own line, and the messages posted on each sign are displayed as horizontal bars. Each bar is positioned on the temporal axis such that the left end of the bar refers to the posting time and the right end refers to when the message was removed. A rendering of the sign's appearance is displayed when users roll the mouse over each message.
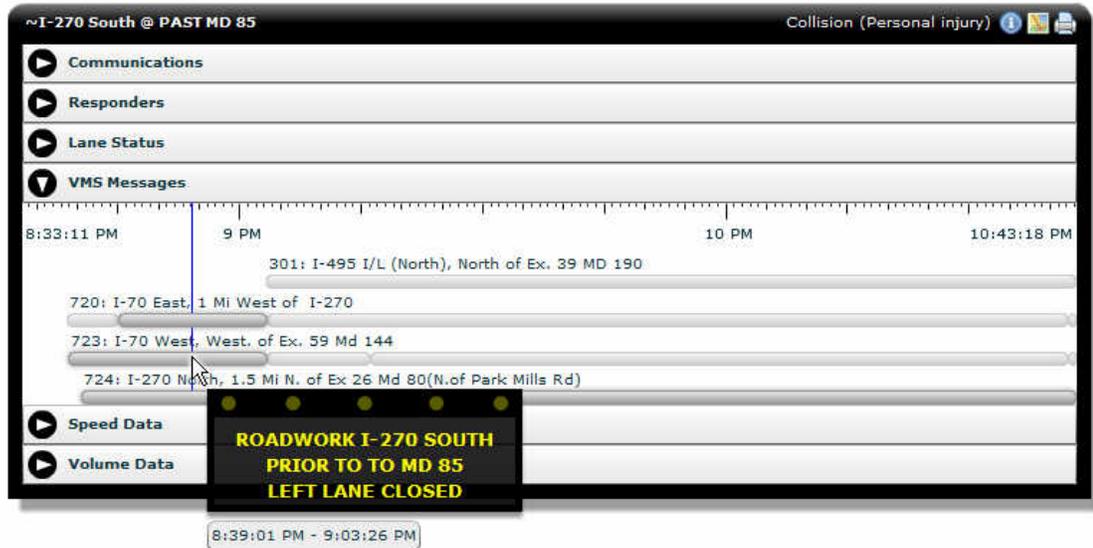
Figure 6: Variable Message Signs rendered in TIME

Knowing the location of each VMS's is useful when determining the appropriate message to post. Using the geospatial coordinates associated with each sign, TIME is able to place the signs and the location of the incident (represented as a red bull's-eye) on an interactive map (Figure 7). The map appears when a user clicks the map icon in the upper-right corner of the application. Like many common mapping tools, TIME allows users to pan and zoom in to a desired location. Clicking on an icon representing a variable message sign will display a miniature timeline for that sign's posted messages.
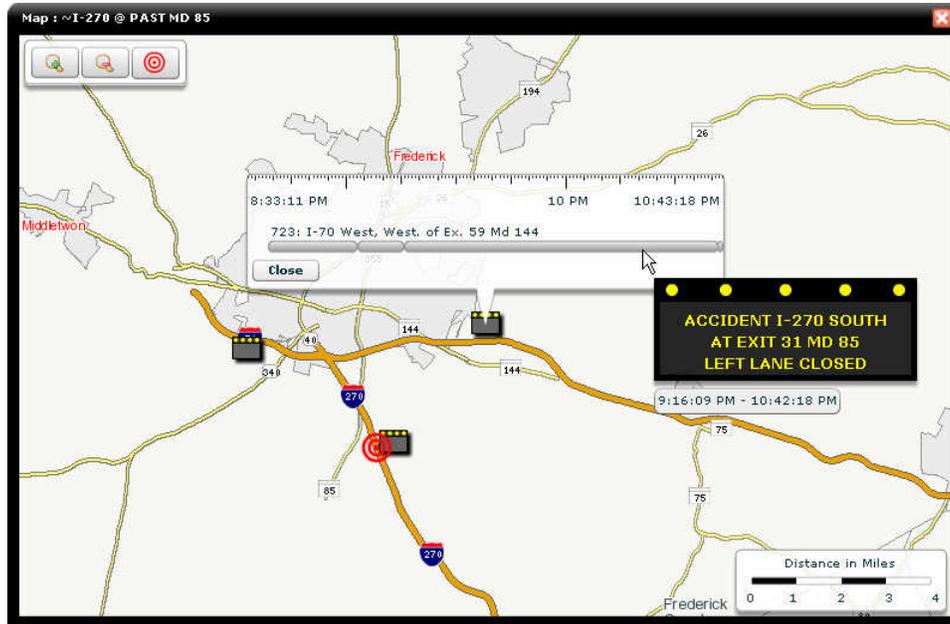
Figure 7: The TIME map with a VMS selected

Responders

The visualization TIME uses for responders (Figure 8) combines the techniques used to display VMS and communications. The notification time of each responder is displayed as a diamond-shaped marker, and the timeframe that the unit was on-scene is represented as a horizontal bar. Tooltips describing each responder event (e.g. "Local police arrived at 9:30:15 pm and departed at 9:50:12 pm") are provided to give users details when needed, and color is used to differentiate the various agencies involved in the incident.

Lane status

Only the current lane status is visible in most existing incident management systems. This is problematic when doing after-action review because the progression of events is impossible to deduce. The visualization used by TIME allows users to view the full history of the incident's lane status. Rather than displaying lane closures in the same style as the interval data in

the other TIME visualizations, a rendering of the highway is used to provide additional spatial information (Figure 8). Lanes are rendered in the same relative locations that they are found on the road; two lanes that are adjacent to each other in the physical world are juxtaposed in the visualization. Time periods when a given lane is opened or closed are colored black or red, respectively. As with the other visualizations, tool tips detailing the duration of lane closures and openings can be revealed by rolling the mouse over the desired segment.
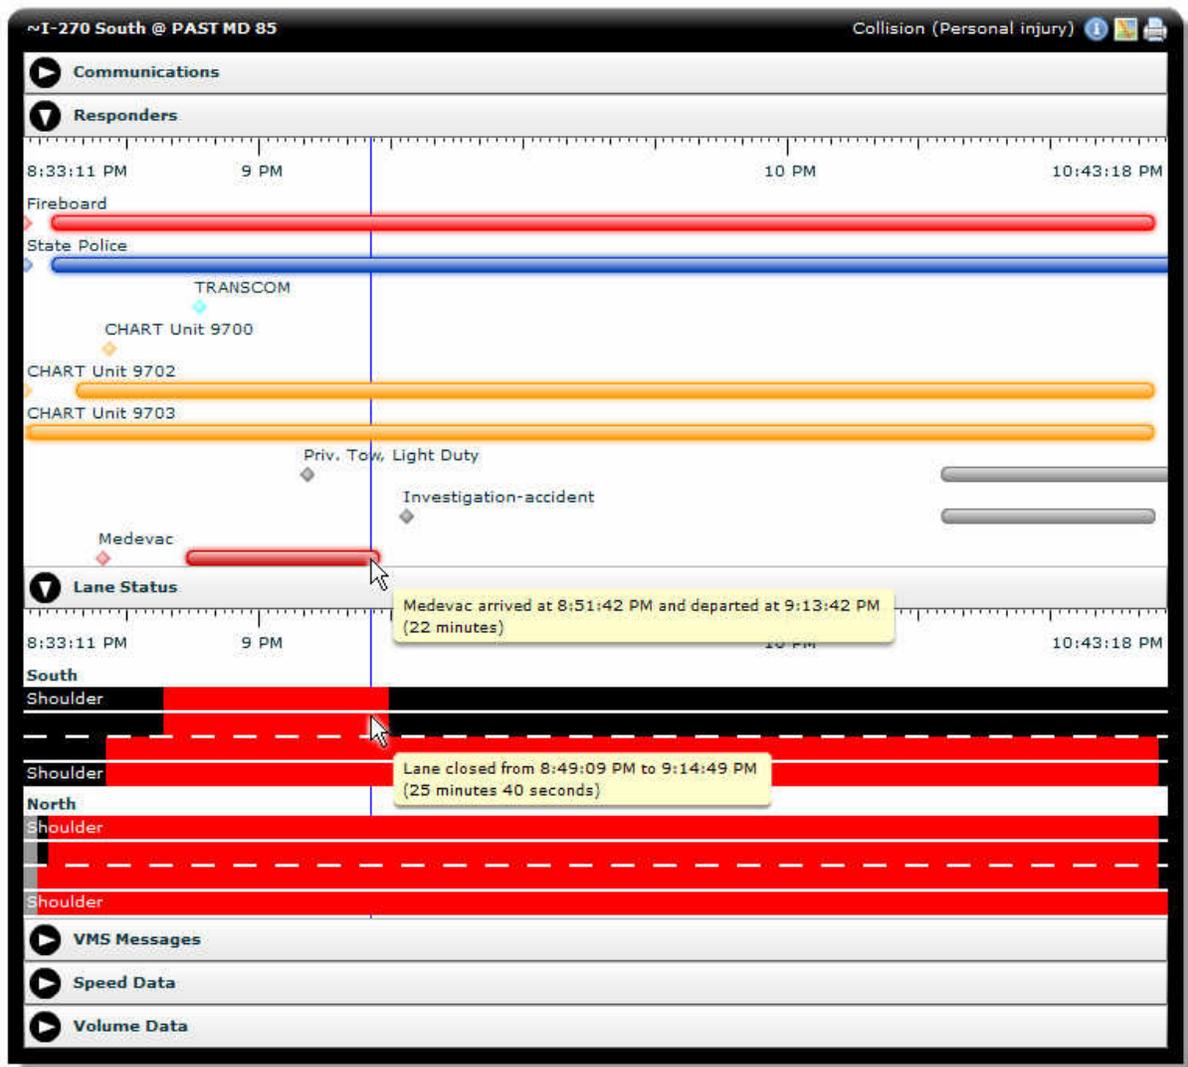

Figure 8: The responders and lane status visualizations

Notice in Figure 8 that a southbound lanes and a shoulder were both closed moments

before a medevac helicopter arrived on scene, and the lanes were reopened as soon as the helicopter departed. This indicates that these lanes were likely closed to provide the helicopter with a place to land. Deducing this series of events is straightforward with TIME.

Speed and Volume Data Graphs

Detector readings provide incident managers with a way to estimate an accident's impact on nearby traffic. As the queue builds up behind the accident scene, detectors upstream of the accident will report lower speed readings. To display this data while maintaining consistency with the other temporal visualizations, TIME converts a traditional timeline visualization into a 2D temporal-spatial plot (Figure 9). The readings for each detector are plotted as pixels horizontally in accordance with when their readings were taken and vertically with the location of the detector. As speeds range from low to high the color varies between red and green. The inverse relationship is used when graphing volume as a lower number of cars on the road is usually more favorable to traffic. Bilinear interpolation is used to smoothly fill in the areas between detectors (vertically) and between readings (horizontally). The vertical axis is labeled with the mile-markers values, and an arrow indicates the direction of traffic flow. A horizontal line is plotted across the entire graph indicating the location of the incident. It is not uncommon to see a patch of red below the line followed by a patch of green above the line. This pattern is indicative of travelers slowing down as they negotiate their vehicles around the incident scene and speeding up immediately after they've passed the location. Rolling the mouse over the graph displays the mile marker, time, and the value of the reading under the cursor.
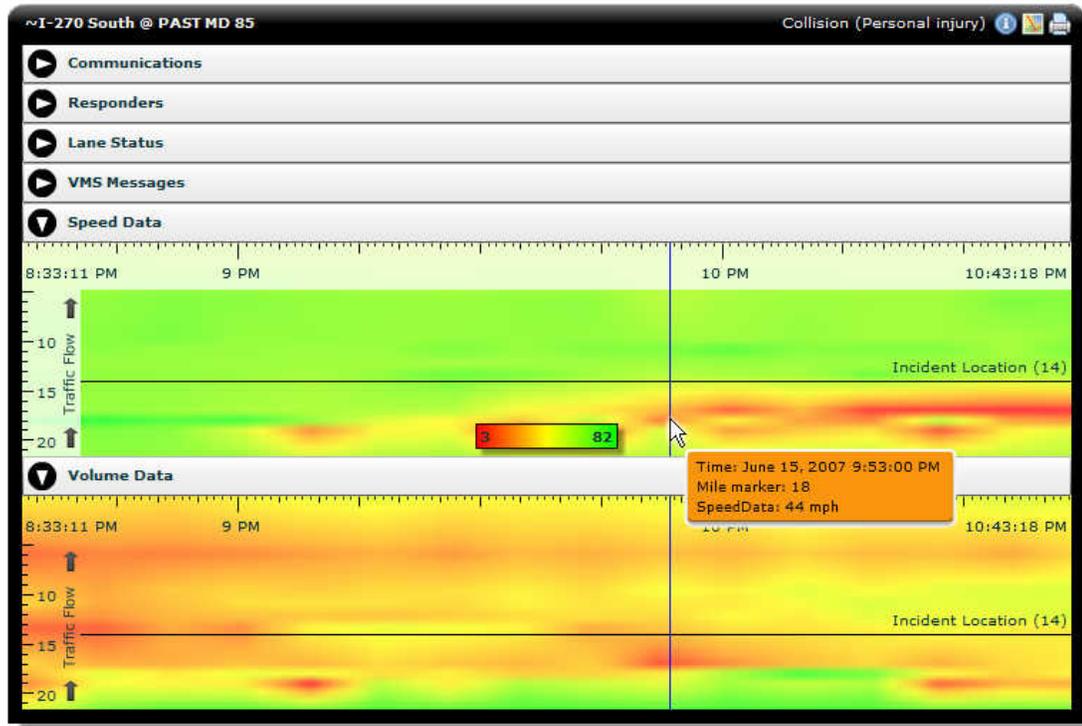
Figure 9: The detector data graphs

None of the incident management systems surveyed provide a way to visualize historical detector data along an entire corridor. In order to view these plots users would have to query the database for the relevant data and enter it into a graphing application. TIME's detector data graphs update automatically as new data becomes available, providing incident managers with a reliable gauge of how an accident affects traffic through a set of data that was previously unavailable to them.

IMPLEMENTATION DETAILS

As incident records are added and modified by the local DOT's the relevant data is pushed into a PostgreSQL database. TIME's server tier is written in ColdFusion and the front-end is written in Flex. The system requires that Adobe Flash Player 9 or higher be installed on the client machine.

TIME's usefulness is entirely dependent upon the usefulness of the data entered by incident managers. Occasionally the data will contain errors or omissions, and the system accounts for these problems on a case by case basis in the most intelligent way it can. For example, operators may indicate that an agency's responder arrived at an incident location without indicating that that agency was ever notified about the incident. In this case the responder would be rendered with an bar indicating their on-scene time and without a notification diamond.

As of this writing there are over 370,000 recorded incidents dating back to June of 2002 in the database and more than 100 more incidents are being logged on a daily basis. Every incident in the database conforms to a format which TIME understands, so even though only a subset of the incidents in the database have ever actually been loaded into TIME (due to human limitations), the system is capable of rendering all of them. It is unlikely that a client machine will have difficulty rendering long running incidents or incidents that have had a lot of activity because the majority of the necessary computation is done on the server. Aggregating the data for the temporal-spatial plots takes significantly longer than aggregating the rest of the data shown in TIME, so the two algorithms are run in parallel rather than sequentially.

Because incident managers are often responsible for several incidents at a time, it is useful for them to view multiple incident timelines simultaneously. To simulate this use case, multiple instances of TIME were opened on a 2 GHz, dual core core machine with 1 GB of RAM. The machine was able to handle thirteen instances before it became too cumbersome to work with. The incidents loaded into the system during the test were randomly selected from all open incidents in the system. Open incidents are significantly more taxing on the client than closed incidents because the system must poll the server every few seconds to see if new data is available. The results of this test are reassuring as incident managers are rarely responsible for more than four or five incidents at a time.

TIME has been available to users for nearly two years, though not all features were present in the original version. Features have been added or modified in accordance with user feedback, and new data sources have been added as they have become available. The repository has seen more than 400 revisions.

USER EVALUATION

TIME is currently being used at VDOT and CHART as part of a regional transportation coordination effort in the D.C. metropolitan region. Although a formal study of the benefits of TIME has not yet been conducted, users' reactions provide some qualitative information about the usefulness of the system. The website that provides access to TIME currently has over 120 registered users. Most of those users were introduced to the system in a similar manner, but this process was by no means standardized. The most common method of introduction involves an in-person tour of the website. During the tour users would be shown how to access the system and how to load both live and archived incident data. While the live incidents shown to each user varied, the archived incidents were hand-picked to provide the most all-encompassing introduction to TIME. These preselected archived incidents have data that lend themselves nicely to an introduction to the system. Many of these users are involved in incident management in some capacity and now work with TIME as a part of their daily workflow.

On average, TIME is accessed 712 times each month by 48 unique users. Only a small subset of the incidents in the database are ever loaded into TIME. Of those incidents, each is viewed by an average of 1.5 users. Each incident's type seems to suggest how many users will load it into TIME. Multi-vehicle accidents and fatalities receive far more views than disabled vehicle towings. This is most likely due to the fact that these large scale incidents require much more attention and coordination from incident managers and garner more interest from other users of the system who may not necessarily be involved in the management process. Although the vast majority of incidents are only viewed by a given user a single time, there are instances where a user will load the same incident between 8 and 10 times. This indicates that users are not only

using TIME as a system awareness tool, but they're also using it as a reference.

Users have commented that the tool is useful for "correlating the multiple types of events that occur during an incident" and detecting when things are not occurring as they should at any given point in time. In addition to diagnosing large scale problems with incident management, users in supervisory positions have commented that TIME allows them to keep tabs on their staff. One user commented that when "the timeline doesn't make sense, it tells me that someone in the control room needs retraining [on the data entry software.]"

The speed at which information can be gleaned from TIME has also received positive comments. Supervisors have expressed satisfaction with TIME as it allows them to quickly review how their operations staffs are handling any given incident. If something is going wrong in the incident management process, a supervisor is able to quickly see the problem and intervene if needed.

For the most part users have favorable comments about how data is presented by TIME. Some of the most frequent requests from users have been to simply visualize other fields in the dataset in the same format that is already used. These additional fields include CCTV video feeds of accident locations, radio communications, and notifications made automatically by the incident management systems themselves. TIME is constantly undergoing development, and these fields will likely be added to the system in some capacity in the future.

The tool has also been successfully used for after-action reviews and by analysts who are evaluating standard operating procedures for incident management centers. Researchers have found TIME to be particularly useful in analyzing both lane closure effects on queues and first-responders effects on lane closures.

FUTURE WORK

Although incident managers can benefit from using TIME in parallel with their current incident management systems, it would be interesting to extend TIME so it can act as a data entry device. A region immediately to the right of the timeline could be made available for making

edits. In the temporal language of the timeline, the far right represents current or pending events. Users could drag responder names or variable message signs into the region, adding them to the log. Through some mechanism, users could even specify events that they would like to occur at some point in the future and the system could remind them to initiate those events when the time comes.

Incident managers are often responsible for large sections of highways that may involve multiple incidents at any given point in time. It would be useful for the managers to be able to view multiple TIME's simultaneously, either side-by-side or in some other fashion that ties the relative locations of the incidents into the visualization.

Similarly, incidents can be related to one another; a rubber-necker gawking at an accident scene might rear-end the car in front of him, causing a secondary accident. Given the volume of secondary incidents logged by incident management systems, it is worthwhile to explore visualizations that combine the data from related incidents. A technique would need to be developed to indicate that these primary/secondary incident pairs share resources because the same group of responders and variable message signs will often be involved in two related incidents.

Individual incidents are often edited by more than one incident manager, so it is difficult to determine accountability for the actions taken during a single incident. A history flow-like interface could provide a easy way for supervisors to oversee the edits of their staff members.

Techniques for handling additional data related to traffic incidents could be incorporated into the system as well. Many incident records include photos, video, radio, and telephone communications, and it is reasonable to assume that these data elements could be incorporated into TIME. These elements along with news clips or articles related to a traffic incident could be valuable additions to TIME especially for after-action reviews.

CONCLUSIONS

TIME is a compact, interactive visualization of the temporal and spatial data associated

with traffic incident management logs. All information is displayed in a compact overview, improving upon previous visualizations which often span five or more separate screens. A variety of techniques are used to convey the usage of information in the dataset. Each visualization is displayed beneath a timescale control representing the duration of the incident. The timescale serves to unify the various categories of data. An interactive map of the incident location allows users to garner spatial significance while still providing temporal information.

The visualizations in TIME make more information readily visible than existing incident management systems, reducing the chance of misinterpreting or overlooking data and expediting after-action analysis. Because incident managers are able to view more data, they are able to make faster, more informed decisions. This in turn leads to quicker incident clearance times and increased transportation system reliability.

ACKNOWLEDGMENTS

REFERENCES

1. Souleyrette R. "US Road Assessment Program," Center for Transportation Research and Education at Iowa State University. May 2006.

2. FEMA, First Responder Training, https://www.firstrespondertraining.gov/home.do , Date accessed: December 22, 2008.

3. I-95 Corridor Coalition, "3D, Distributed, Interactive Transportation Management & First Responder Training System," http://www.i95coalition.net/i95/Training/IncidentMgmtVirtualTraining/tabid/189/Default.aspx , Date accessed : December 22, 2008.

4. Aigner W. "Visual Methods for Analyzing Time-Oriented Data," IEEE Transactions on Visualization and Computer Graphics. Vol. 14. 2008. pp 47-59.

5. Aigner W., Miksch S., Thurnher, Biffl S., "Planning Lines: Novel Glyphs for Representing Temporal Uncertainties and Their Evaluation," Proc. Ninth Int'l Conf. Information Visualization (IV '05), 2005.

6. Havre S., Hetzler B., Nowell L. "ThemeRiver: Visualizing Theme Changes over Time", INFOVIS '00: Proceedings of the IEEE Symposium on Information Vizualization 2000.

7. Buono P., Plaisant, C., Simeone, A., Aris, A., Shneiderman, B., Shmueli, G., Jank, W. "Similarity-Based Forecasting with Simultaneous Previews: A River Plot Interface for Time Series Forecasting", Proc. of the 11th International Conference Information Visualization (IV '07), 2007, 191-196.

8. David François Huynh, "SIMILE | Timeline" http://code.google.com/p/simile-widgets/ Date accessed: December 7, 2008.

9. Taowei David Wang, Catherine Plaisant, Alex Quinn, Roman Stanchak, Ben Shneiderman, and Shawn Murphy. "Aligning Temporal Data by Sentinel Events: Discovering Patterns in Electronic Health Records", SIGCHI Conference on Human Factors in Computing Systems, 2008.

10. Viégas, F., Watternberg, M., and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. Proceedings of CHI'04, (Vienna, Austria), 2004, 575-582.

11. BBC, "British History Timeline"
    http://www.bbc.co.uk/history/interactive/timelines/british/index.shtml ,Date accessed:
    December 7, 2008.

12. Popular Mechanics "Space History - First 50 Years - Timeline",
    http://www.popularmechanics.com/science/air_space/4221976.html , Date accessed:
    December 7, 2008.

13. Google Calendar, http://www.google.com/googlecalendar/overview.html , Date accessed:
    December 7, 2008.

14. Microsoft Outlook, http://office.microsoft.com/en-us/outlook/default.aspx , Date
    accessed: December 7, 2008.

15. Mozilla Sunbird, http://www.mozilla.org/projects/calendar/sunbird/ , Date accessed:
    December 7, 2008.

16. Luz, S., Masoodian, M. "Visualisation of Parallel Data Streams with Temporal Mosaics,"
    11th International Conference Information Visualization. 2007.

17. Plaisant, C., Milash, B., Rose, A., Widoff, S., Shneiderman, B., "LifeLines: Visualizing
    Personal Histories", in Proc. of CHI 96, ACM, New York.

18. ESRI, "ArcGIS: The Complete Enterprise System",
    http://www.esri.com/software/arcgis/index.html , Date accessed: December 7, 2008.

19. Google Earth, http://earth.google.com/ , Date accessed: December 7, 2008.

20. MacEachren, A., X. Dai, F. Hardisty, D. Guo, & G. Lengerich, "Exploring High-D Spaces
    with Multiform Matricies and Small Multiples," Proceedings of the International
    Symposium on Information Visualization, 2003.

21. Kapler, T. and W. Wright, "GeoTime Information Visualization", Information
    Visualization Journal, Palgrave Macmillan, 4,(2), Summer 2005, 136-146.

22. Andrienko, G., Andrienko, N., "Spatio-temporal Aggregation for Visual Analysis of
    Movements," IEEE Symposium on Visual Analytics Science and Technology, p. 51-58.
    2008.

23. Andrienko, G., Andrienko, N., "Visual Exploration of the Spatial Distribution of

Temporal Behaviors," Proceedings of the Ninth International Conference on Information Visualisation. 2005.