

ANALYZING INCIDENT MANAGEMENT EVENT SEQUENCES WITH INTERACTIVE VISUALIZATION

John Alexis Guerra-Gómez
Human-Computer Interaction Lab &
Department of Computer Science
2105 Hornbake Building
University of Maryland
College Park, Maryland 20742
+1 740 591 0435
jguerrag@cs.umd.edu

Krist Wongsuphasawat
Human-Computer Interaction Lab &
Department of Computer Science
University of Maryland
College Park, Maryland 20742
kristw@cs.umd.edu

Taowei David Wang
Human-Computer Interaction Lab &
Department of Computer Science
University of Maryland
College Park, Maryland 20742
tw7@cs.umd.edu

Michael L. Pack
Center for Advanced Transportation Technology
University of Maryland
College Park, Maryland 20742
+1 301 405 0722
packml@umd.edu

Catherine Plaisant
Human-Computer Interaction Lab &
Institute for Advanced Computer Studies
University of Maryland
College Park, Maryland 20742
+1 301 405 2768
plaisant@cs.umd.edu

Number of words in abstract, text and references: 4,864. Number of figures: 8
Total effective number of words = $4,864 + (8 * 250) = 6,864$

ABSTRACT

While traditional safety and incident analysis has mostly focused on incident attributes data, such as the location and time of the incident, there are other aspects in incident response that are temporal in nature and are more difficult to analyze. We describe a visual analytics tool for temporal data exploration, called LifeFlow, used for the analysis of incident response data. LifeFlow provides user-controlled overviews of event sequences (e.g. notification, arrival, clearance etc). It allows analysts to interactively explore temporal patterns, find anomalies in sequences and compare management practices. This type of analysis can potentially lead to process improvements and save human lives. We used NCHRP traffic incident data with more than 200,000 incidents are reported by 8 different agencies in a period of about 28 months. Our experience suggest that even non expert analysts can spot many anomalies in the data using the LifeFlow overviews, and are able to rapidly ask many questions and find differences between agencies.

INTRODUCTION

Vehicle crashes remain the leading cause of death for people between the ages of four and thirty-four. In 2008, approximately 6-million traffic accidents occurred in the United States. This resulted in nearly 40,000 deaths, 2.5 million injuries, and losses estimated at \$237 billion. Given these statistics (1,2), reducing the impacts of crashes and improving safety on the nation's highways has become a significant issue for policymakers and transportation engineers. To combat this problem, transportation systems are monitored at an unprecedented scope which results in tremendously detailed traffic, road weather, and comprehensive incident databases. While the transportation community emphasizes developing standards for archiving and transmitting raw incident data, little effort has been made to design appropriate visual analytics tools to explore the data, extract meaningful knowledge, and represent results. We believe that interactive visualization tools can help transportation managers and government agencies better analyze these large multivariate datasets (3).

While traditional safety and incident analysis has mostly focused on incident attributes data (such as the location and time of the incident or type of vehicle involved, there are a significant number of other variables in incident management that that are temporal in nature and are more difficult to analyze. These variables may include what occurred prior to the incident, e.g. a driver consumed alcohol, amount of time spent driving, length of time since water/chemical spill, or amount of time since sunset. Once the incident occurs there are other important variables such as when responders and equipment arrive on scene, the length of time that lanes are blocked or closed, the time it takes for a queue to reach a critical point, etc.

Specifying complex temporal queries in traditional database query languages, such as SQL is extremely difficult even for computer scientists specializing in such analysis. Furthermore, proper analysis with traditional SQL requires the analyst to have a clear understanding of the questions to be asked, which often results in many important data relationships being overlooked. Scientific research has made progress in representing temporal abstractions and executing complex temporal queries but there is very little research that focuses

on making it easy for users to specify the queries and to examine the results visually. Some research systems provide temporal access languages to support limited visual queries from end-users but these suffer the same accessibility difficulties of SQL, namely that the languages are not known by nor are they welcoming to practitioners, or they require an understanding of the underlying database structure. Simply stated, these query languages are too obtuse to use.

In this paper we apply and extend LifeFlow (4), a new visual analytics tool for temporal data exploration, to the analysis of incident data. LifeFlow provides user-controlled overviews of event sequences. It allows analysts to visually explore temporal patterns in the data, find anomalies in event sequences and compare management practices based on incident attributes (*e.g.* incident type or agency type.)

We first summarize related work and describe our case study data, then describe the LifeFlow visualization and user interface. Finally we illustrate the use of the tool with our case study data, and provide example of early findings. We conclude with feedback from transportation experts and possible future work directions.

PREVIOUS WORK

While static visual displays of information have been used to summarize and present information for a long time (5), only recently have interactive systems become available. In the field of incident traffic analysis, analysis remains often done using traditional Excel spreadsheets and SQL queries (3). However, as transportation data have grown larger and more complex, visual analytics tools are also starting to emerge.

Visual analytics tools have been used to analyze the timing and location of incidents, *e.g.* (6). ICE (Incident Cluster Explorer) combines interactive maps, histograms, two-dimensional plots and parallel coordinates plots and allows users to interact simultaneously see relationships between the tightly coupled visualizations to analyze patterns in multidimensional incidents incident data (7,8).

Using simulations in the backend interactive visualizations have been used for training, *e.g.* (3) describe a system to provide designers and operators with an appreciation of the impact of highway incidents on traffic delays on arterials. The ambitious Virtual Incident Management (3), a 3D massively multiplayer virtual environment, shows great promise for training traffic control personnel, emergency management personnel and first responders. These environments can help trainees obtain a better understanding of the traffic incidents in a safe environment, while generating large amounts of data that can be analyzed later or in real time.

To represent incident response data, tools that focus on categorical temporal data should be considered. An example is Lifelines (10) which presents a single screen overview of a medical patient record. This concept was applied to the display of incident data (11). Another interactive visualization, Lifelines2 (12) allowed users to explore collections of records instead of a single record and to search for specific temporal patterns (*e.g.* find all records including event A followed by event B within 2 hours). This work demonstrated that interactive visual exploration can beneficially complement query formulation. The features in Lifelines2 were

found to be useful by medical researchers for tasks such as identifying patients for clinical trials, establishing quality assurance measures, testing alarm conditions, etc. Similarly, Lifelines2 can be used to find all incidents with specific known patterns, *e.g.* all incidents for which lane clearance occurred more than 5 hours after a lane closure. One limitation of Lifelines2 is that it does not provide an overview of all the records at once, making difficult to see the most common temporal sequences or to spot unusual or anomalous ones -- tasks that our new work on LifeFlow addresses.

Interfaces have also been developed to find patterns similar to a template or a reference record. Similan (13) allows sequences search by similarities while retaining LifeLines2's main visualization, and explores different ways of specifying and performing the temporal search using a customizable similarity measure. Every record is ranked by similarity to the reference one, revealing variants of the pattern. However, Similan shares Lifelines2 main limitation: the need to know beforehand the sequences to be searched, and the absence of a temporal sequence overview technique.

Finally, the LifeFlow display itself was inspired by a hierarchical visualization called Icicle Tree (14) which uses a space filling technique to summarize overviews of categorical data in a fixed space area. Like Treemaps (15) this technique is useful to convey characteristics of a hierarchy, such as the distribution of nodes in the branch of the hierarchy. LifeFlow uses a drawing algorithm similar to the Icicle Tree but adds a time dimension in the *x* axis while keeping the proportions of event sequences along the *y* axis.

DATASET USED IN THIS STUDY

To guide our design and implementation of LifeFlow for the analysis of incident data, we used a National Cooperative Highway Research Program (NCHRP) dataset that includes 203,214 traffic incidents from 8 agencies.

In incident includes two attributes: the agency that processed it (represented with a letter from A to H for anonymity) and the type of incident (*e.g.* “*Disabled Vehicle*”, “*Fatal Accident*”, “*Multiple Collision*”, etc.) The values for incident types are textual and not standardized across the agencies, *e.g.* “*Disabled*” versus “*Disabled Vehicle*”.

In addition, each incident has a sequence of incident management events: “*Incident notification*”, “*Incident arrival*”, “*Incident clearance*”, “*Incident cleared*”, “*Lane clearance*”, “*Return to normal*”. The meanings of the events can also vary from agency to agency.

A typical incident sequence includes:

- “*Incident notification*” marking when the agency is first notified of the incident
- “*Incident Arrival*” marking when the emergency team arrive the scene.
- “*Lane Clearance*” marking when the lanes are opened, but the incident scene may not being completely cleared yet.

- “*Incident cleared*”, “*Incident clearance*”, and “*Return to normal*” all denote the end of incidents. Later on in the analysis, we aggregated all three into the new event type “*Return to normal (aggregated)*.”

A typical sequence of event starts with “*Incident Notification*” and finish with “*Return to normal*”, with the possibility of having “*Incident Arrival*” and “*Lance Clearance*” between them.

LIFEFLOW DISPLAY

When data is first loaded into LifeFlow, the software first analyzes it to find all possible sequences of events across all incident records (even those the analysts may not expect to find.) It then aggregates all incident records with the same sequence into a tree of sequences, and creates a visualization of this tree. For example, if three incidents have the same sequence of four events, these three incidents are aggregated together (Figure 1). A color-coded vertical bar represents each event on the common sequence. In this example, there are four events, so four vertical bars are used. The distance between the bars indicates the mean (average) time between events in the set of incidents with the same sequence. The height of the aggregation represents the number of records or incidents. Figure 2 shows an example of the process followed to display the data from its original representation to the LifeFlow display.

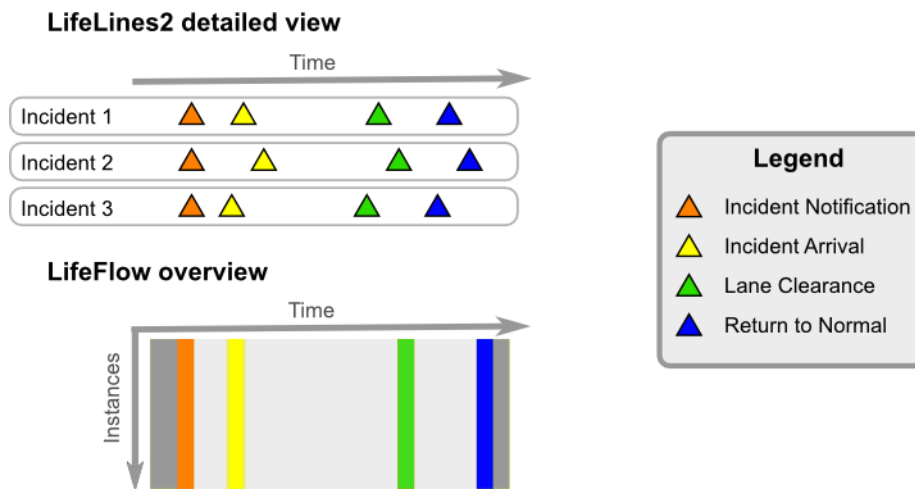


FIGURE 1 A single sequence bar aggregates multiple incidents with similar sequence of events: “*Incident Notification*” (orange), “*Incident Arrival*” (yellow), “*Lane Clearance*” (green), and “*Return to Normal (Aggregated)*” (blue). The distance between the bars is proportional to the average time between events. The height of the bar is proportional to the number of incidents in the aggregation.

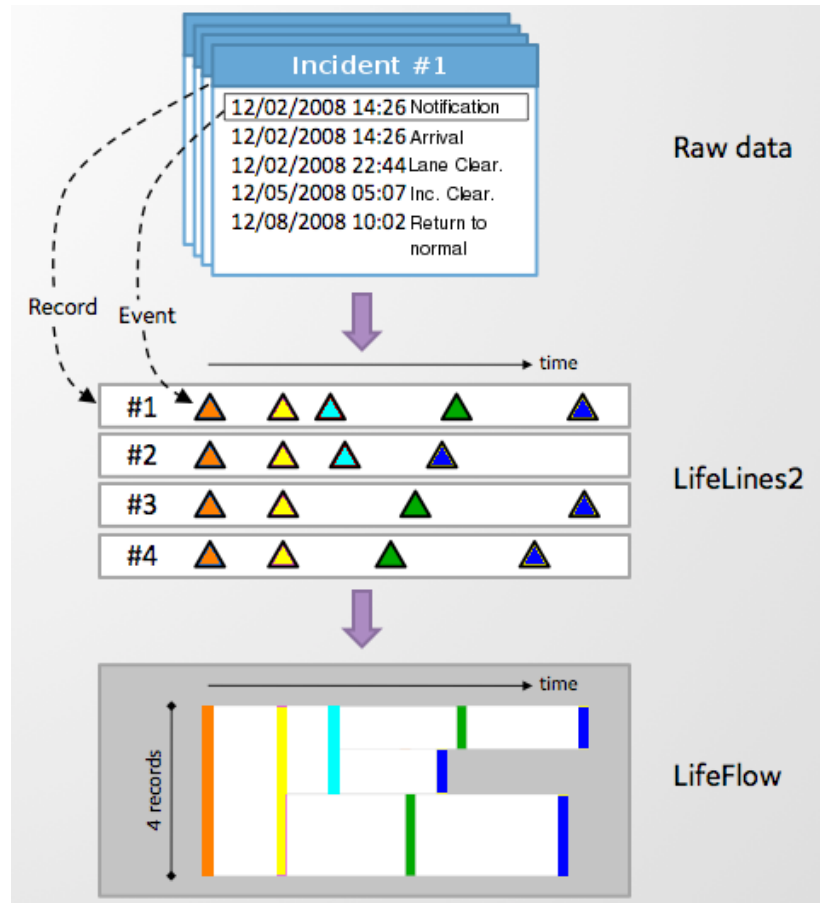


FIGURE 2: Constructing a LifeFlow display: The raw table data in the incident log can be displayed on LifeLines 2 as a set of triangles colored by incident type. Then all incidents exhibiting the same sequence are aggregated in sequence bars, and the final LifeFlow display combines all the sequence bars so all the sequences in the data can be represented at once.

Figure 3 shows an example with 5106 real incidents. The large top sequence bar tells us that most of those incidents follow a normal sequence (a) “*Incident Notification*”, “*Incident Arrival*”, “*Lane Clearance*” and “*Return to Normal*”, with an average of 22 minutes and 41 seconds between notification and lane clearance, and another 9 minutes and 14 seconds before return to normal. In this example only a few incidents have no line clearance event, but many incidents have a sequence where the clearance is recorded as occurring before arrival on the scene (552 incidents). However the display also reveals many other several unexpected sequences (b) where, for instance, the “*Lane Clearance*” occurred before the “*Incident Notification*”, also shown in the detailed view of Figure 4. Those anomalies have shorter average time to clearance.




FIGURE 3 LifeFlow display showing 5,106 incidents. Most of those incidents follow a normal sequence (a) “Incident Notification”, “Incident Arrival”, “Lane Clearance” and “Return to Normal”, with an average of *about 23 minutes* before lane clearance and *about 32 minutes* before return to normal.

USER INTERFACE

LifeFlow can summarize of thousands of records in one screen, but also enables the analysts to interact with the data in order to answer questions they have. The following is a description of the operations that users can perform on LifeFlow:

1. Access to Details.
2. When users move the mouse cursor over a sequence bar, a tooltip describing details of the bar appears. The detail includes basic statistics of the time to the previous event and to the first event, etc. An example of popup dialog can be seen in Figure 4.
3. When users select a sequence bar, a detailed view allows the user to look at individual records. A window opens on the right side with a LifeLines2 timeline display. Both views

are tightly coupled, so when the user selects a group of incidents in the overview, those same incidents appear in the detailed view for further analysis. See Figure 4 on the right.

4. **Split Display by Attributes.** Since the input data may also include non-temporal attributes for each of the incidents (*“Incident Type”* and *“Agency”*) users can split the display accordingly to group all incidents by those attributes. To achieve this, users select the type of attribute they want to use, for instance *“Agency”*. This causes LifeFlow to aggregate incidents by their agency type first. The result is that there is a sub-tree for each agency in the display, facilitating cross-agency comparison. The agency label is displayed as a big box on the left of each sub-tree, allowing easy identification and selection of each individual agency. Figure 5 shows an overview of all the different sequences for all the agencies. The agencies identifiers (letters A to H) can be seen in white boxes on the left of the overview.
5. **Show/Hide event types.** Using the legend on the left side of the screen users can show or ignore certain events using the check boxes. For example by removing the *lane clearance* events the display is regenerated to combine all the records with similar sequences whether they have lane clearance events or not. This allows users to aggregate the data as needed to compute the mean times between events on the new aggregates.
6. **Selection and Filter.** By clicking on any vertical colored bar on the overview, users can select groups of incidents (*e.g.* select all the incidents that begin with *“Incident Notification”*). After selecting instances users can filter instances by choosing to keep only the selection and remove everything else, or the opposite. This is an easy way to remove anomalous data that should be eliminated from further analysis.
7. **Align.** By default all the events are displayed without alignment that is presenting the sequence bars in order of occurrence in the overview and presenting the events in an absolute timeline in the details. However sometimes users need to analyze more detailed what happens before and after a specific event. For this users can select to align the visualizations by any of the incident events. By doing so the detailed view presents a relative timeline for events showing the time difference from the rest of events to the aligned event (as shown in FIGURE 4 on the right). On the other hand, when alignment is made, the overview is divided in two sections one for the events before the alignment and one for the events after it.
8. **Select Common Sequences Across Attribute Values.** After displaying the non-temporal attributes, users can also select a sequence across different attribute values. For instance, selecting the sequence *“Incident Notification”, “Incident Arrival”, “Return to Normal”* in an attribute sub tree will select the same path in all the different agency sub trees. After making these selections the users can likewise apply filter.
9. **Measure Time Range.** LifeFlow provides a measuring tool that gives the time distance between two nodes of the overview tree (dotted line in Figure 4), allowing users to see the exact difference in the average time to clearance between of agencies A and B, for instance. In Figure 4 the measuring tool shows that in average 29 minutes pass between *“Incident Arrival”* and *“Return to Normal”* in the most common sequence. Finally the

details popup (d) is also shown describing some details about the time gap between the “Lane Clearance” and the “Return to Normal”.

10. **Sort.** To facilitate comparisons, users can sort the layout using a menu option in different ways: 1) according to the number of incidents the bars represent, 2) the longest (in time) sequence bar, or 3) the average length of the bars. Users can order the leaves of the overview tree by these three criteria, and facilitate comparisons. For instance, one can quickly sort agencies by their average time to clear the incidents, or by the number of instances they reported.
11. **Displaying All Sequence Bars with Equal Height.** When data includes a large number of incidents, it can be difficult to spot the aggregates that represent just a small number of incidents (*e.g.* dozens out of thousands). An option displays all the sequences using equal height, which makes it easier to review all possible sequences in the dataset.
12. **Load/Save Data.** The basic format for LifeFlow is tab-delimited text files, where the columns are incident ID, event type, event time stamp, and annotation. Users can also export the results of their explorations.

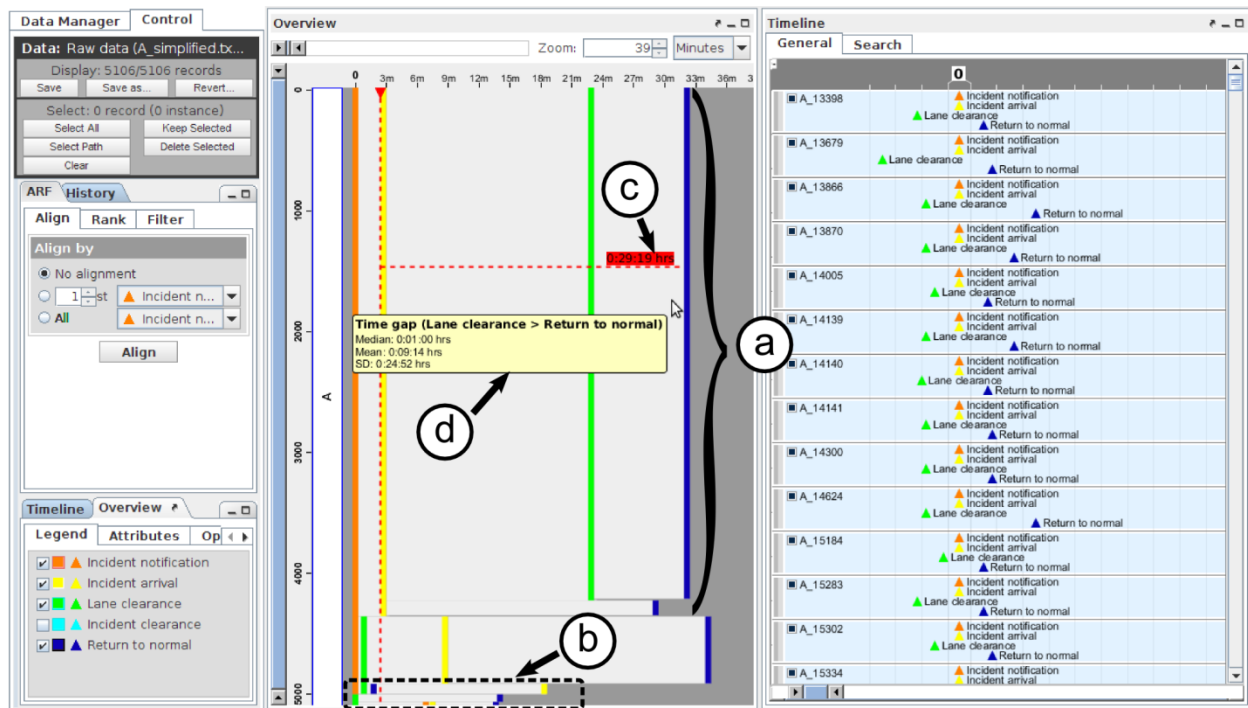


FIGURE 4 LifeFlow Interface. The main LifeFlow windows are divided in three sections: on the left the control panel, on the middle the LifeFlow overview, and on the right the detailed Lifelines2 view. Here we are representing 5106 incidents attended by the Agency A. We can see in the overview that most of those incidents follow a normal sequence (a) “Incident Notification”, “Incident Arrival”, “Lane Clearance” and “Return to Normal”, However there are also some unexpected sequences (b) where, for instance, the “Lane Clearance” occurred before the “Incident Notification”, also shown in the detailed view.

CASE STUDY – RESULTS AND FINDINGS

We used the dataset analysis mostly to guide our design. Nevertheless, we were able to explore the data in a limited way and made several observations. We use those observations here solely to illustrate the use of the tool. Further analysis may be needed to validate our preliminary observations.

Quantifying data quality issues

When the entire dataset of 203,214 incidents was loaded into LifeFlow, only 195,346 out of all incidents exhibited the expected sequence of events. Many variations and anomalies were found.

After loading the dataset in LifeFlow, we immediately noticed that the Agency B contains a path that is more than 110 years long with 6,712 incidents (Figure 5). Investigating deeper into this path, we found that Agency B reported the “*Incident Arrival*” of those incidents as January 1th 1900. Since this date is commonly used as the initial date in computer systems, this suggests that the system the Agency used to register this event might have used it as a default value. Considering this part of the data as corrupted, we remove all incidents in this part from the dataset. While it was easy to spot this problem, such anomalies can often remain undetected, and skew the results of even the simplest of analysis such as calculating the mean time to clearance.

Similarly, we found 48 incidents from Agency D that are about 10 months long, in which the “*Incident Arrival*” occurs before the “*Incident Notification*”. Considering those data corrupted as well, we also removed them.




FIGURE 5 NCHRP dataset LifeFlow overview. The long path in the middle shows that there is a very anomalous group of incidents that are in average 110 years long.

The next thing we noticed from the resulting data was that there were many incidents that lasted exactly 24 hours, which seemed unlikely. Using the Align, Rank and Filter operators LifeFlow inherited from Lifelines2, we found that those 24-hours-long incidents have “*Incident arrival*” events occur in the first hour of the day (e.g. 12:30AM April 10 2009) and “*Incident notification*” events happened in the last hour of the same day (e.g. 11:50PM April 10 2009). This observation seemed to suggest that there were data entry problems with those incidents, indicating that the operator failed to - or was not able to - record the correct date of an event (e.g. 12:30AM April 11 2009 as opposed to 12:30AM April 10 2009). Similar errors were discovered for paths that are about 12 hours long, in which case the errors seem to be problems choosing between AM and PM in the date.

Those anomalies were found quite easily by the computer scientist main developer, who had no experience in transportation data. Finding such errors using traditional tools like SQL or manual analysis can be very difficult and time consuming, and requires experienced analyst who would suspect the existence of such errors.

Ranking the Agencies’ Performance

In this study, we used the time from when the agencies were notified to the final clearance of the incidents as a performance measure.

The time when the agency was notified in each incident can be indicated by the “*Incident Notification*” event. In order to compare the agencies’ performance, we needed to remove all the inconsistent data first. We removed all the incidents that do not start with “*Incident Notification*”, which could be performed easily using the equal height overview feature.

However, these agencies use different event types to indicate the final clearance of incidents (e.g. “*Incident cleared*”, “*Incident clearance*” and “*Return to normal*”). Therefore, we replaced these three final clearance event types with a new event type called “*Return to normal (Aggregated)*”.

After the steps above, the visualization of the data can be seen in Figure 6. Incidents are grouped by agencies. We showed only two event types (“*Incident Notification*” (orange) and “*Return to Normal (Aggregated)*” (blue)), so the horizontal length of each agency’s path represents the average time from incident notification to final clearance, which reflects the performance measure for that agency. We then sorted the agencies according to the length of their paths, resulting in the fastest agency (shortest path) on the top and the slowest agency (longest path) in the bottom.

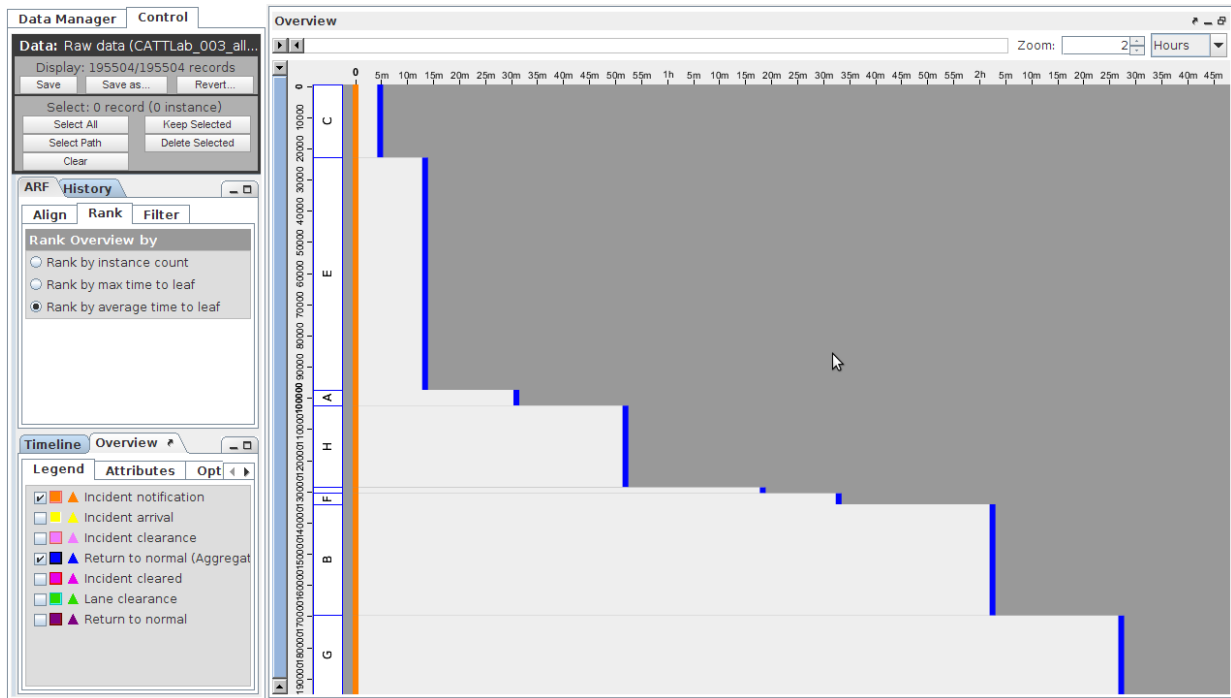


FIGURE 6 Sorted display by a simple measure of agency performance. C seems to be the fastest to clear its incidents, followed by E, A, H, D, F, B and finally G. Access to details will of course be needed to look for clues as to why those differences exist.

From Figure 6 we can see that Agency C was the fastest agency to clear its incidents, taking about 5 minutes in average, while the slowest one was Agency G with an average of about 2 hours 27 minutes.

To investigate deeper into Agency C’s data, we removed the data from other agencies. Then we looked into the different Incident Types reported. We found that most of the incidents that Agency C reported are “*Disabled Vehicles*” which had in average a time to clear of about 1 minute (Figure 7). Looking at the event distribution, we also found that a large number of the incidents reported “*clearance*” immediately after “*Incident Notification*”. This observation made us wonder if there is any explanation for these immediate clearances, and encouraged further analysis.

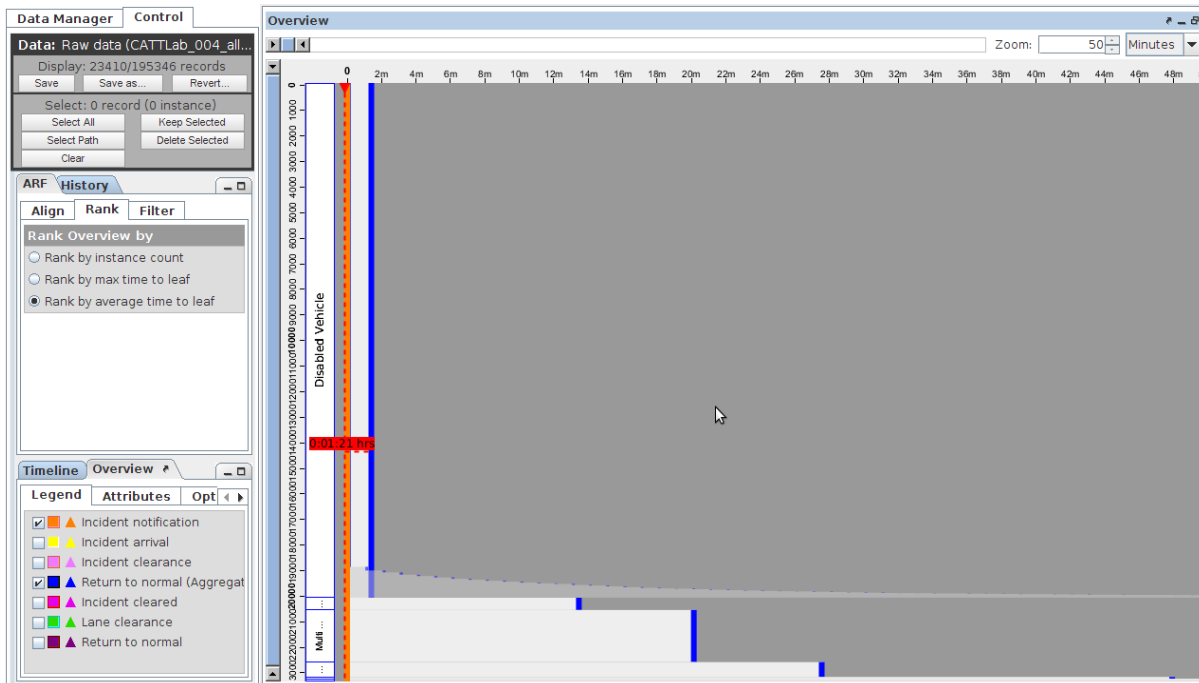


FIGURE 7 Overview of incident for the Agency C by Incident Type. This overview shows that the big majority of incidents covered by this agency are “Disabled Vehicles” with an average return to normal time of 1 minute 21 seconds. Even more the distribution of those events show that most of them were cleared immediately after the “Incident Notification” which is unlike other practices where it takes time to respond to a disabled vehicle notification.

In a similar fashion, we investigated Agency G, which seems to be the slowest agency, in more details. Figure 8 shows the overview of the Incident Types for the Agency G. Agency G classifies their incidents in only two types “Non-ATMS Route Incident” and simply “Incident”. As you can see from the visualization in Figure 8, the group with Incident Type “Incident” has an average length of about 38 minutes, which is a very good time compared to the other agencies, but the group with Incident Type “Non-ATMS Route Incident” took in average 5 hours 14 minutes to clear. So we realized that when using the average time of all incidents from Agency G without considering the incident types, Agency G seems to be slower than other agencies. While in fact, Agency G performs quite well for incidents with Incident Type “Incident”.

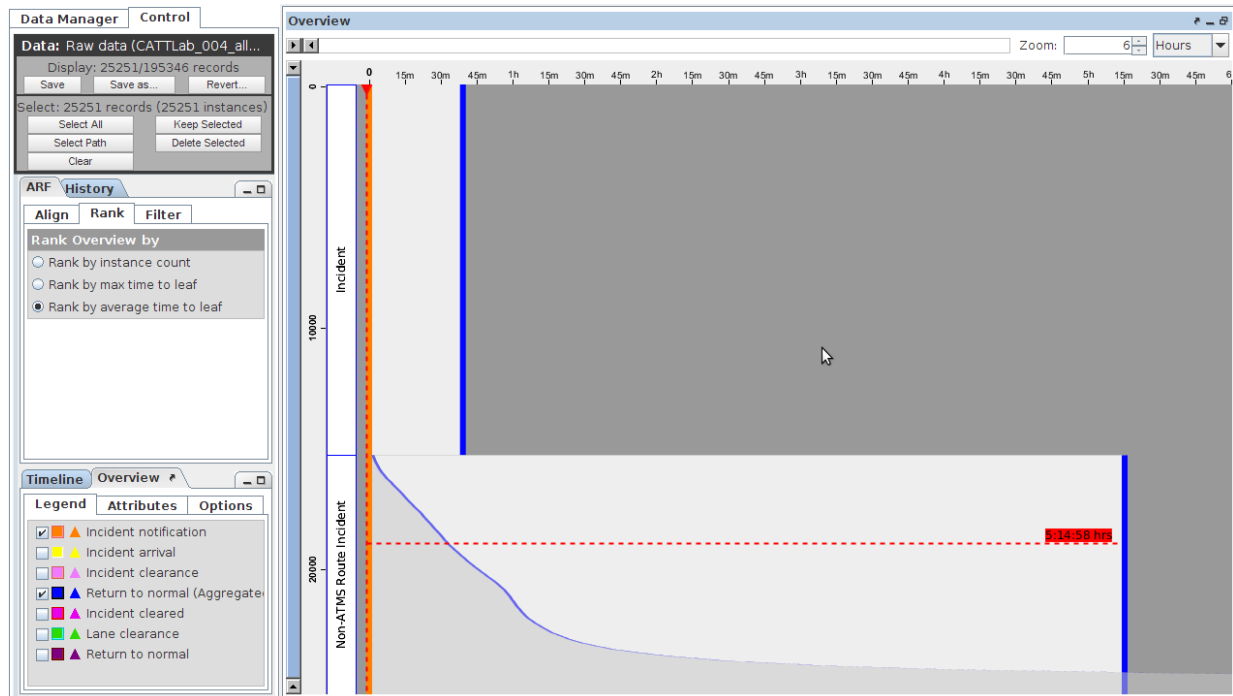


FIGURE 8 Incident types overview for the agency G. This shows that the Agency handled only two types of incidents, and that the “Non-ATMS Route Incident” took them in average 5 hours 14 minutes to clear, while the other Incidents were cleared in fewer than 45 minutes.

DISCUSSION

Although our data analysis in the case study is limited and preliminary, domain experts from the CATT Lab is conducting a more formal analysis of the data. They reviewed our work and stated that they wished LifeFlow was available earlier on when they started their own analysis. They confirmed the existence of anomalies that we had found in the data, and stated that their elimination was non-trivial when using SQL. For example, when the analysts calculated the average time it took an agency to clear its events, they had to run an SQL query that searched for: “The average time difference of all the incidents where the *Return to Normal* is after the *Incident Notification*”. This means that they had to expect this error in the data, and they had to be careful to exclude them from their analysis. However excluding all the possible erroneous sequences in a SQL query would be very difficult. In the end, they needed to review the results of SQL queries to ascertain that there were no longer any errors. Without LifeFlow, this kind of review and identification of unexpected sequences would be almost impossible. Finally, they mentioned that LifeFlow would allow them to ask more questions faster, and probably richer questions about the data. LifeFlow was also able to reveal unexpected sequences that may have been overlooked, but the tool also suggested that their prevalence is limited.

We believe that using LifeFlow can assist analysts explore large datasets, such as the NCHRP traffic incident information, in ways that would be very difficult using traditional tools. This might allow analysts to find richer results in less time. For example, we believe that

LifeFlow can help analysts identify potential factors for the agencies' differing performance results, and can quickly isolate data for further analysis.

Challenges

1. **Normalization:** Agencies report their incident management timing data in inconsistent ways. Difference may stem from differences in data entry practices, terminology variations, or missing data. Policies are needed to decide if, and how, to normalize the data and tools will be needed to manage this normalization process and provide audit trails documenting the transformations. Beside timing information, the same problem occurs with incident types as the terminology is not consistent across agencies or even local areas. This leads to a very large number of incident types in the interface and restricts comparisons. Tools to interactively group and ungroup events types may be valuable.
2. **Scalability:** While our initial prototype can already load 400,000 incidents, interaction is slow with such numbers, until the total number can be filtered down to smaller subset of interest. We will need to refine our algorithms to handle even larger numbers and evaluate their performance, but we are confident that this problem can be addressed.
3. **Interface complexity:** Interfaces with many features and options can become difficult to use. Future improvements will strive for simplicity, carefully making important features directly accessible while relegated rare features to menus and control panels.

CONCLUSION

In this paper we present a new visual analytics tool called LifeFlow and demonstrate how it can be used to explore incident data. Our case study looked at more than 200,000 incident data from 8 transportation agencies. A rapid analysis by non experts reveal many unexpected sequences in the reported data (events dating more than 100 years ago, unusually long paths, possible AM/PM input errors and other anomalous sequences). We used LifeFlow to compare the performance of the agencies in clearing the incidents, and drilled down quickly to explore potential factors that affect performance.

We believe that the analysts can use LifeFlow to ask more, richer questions than using SQL or direct analysis of the data in spreadsheets. LifeFlow allows exploration of the data, revealing not only common sequences, but also unexpected, erroneous sequences that would have ruined any statistical analysis if the errors were not fixed in the first place.

LifeFlow is still under active development, and we have many plans for future expansion of features. By allowing better and faster understanding of the always-growing traffic incident data, LifeFlow may help identify best practices and ultimately help reduce congestion and save human lives.

ACKNOWLEDGEMENTS

We would like to acknowledge the “Center for Integrated Transportation Systems Management, a Tier 1 Transportation Center at the University of Maryland” for partial support of this research; and also the National Institute of Health (Grant RC1CA147489-02) for supporting our research on temporal pattern analysis of patient records which led to the original design of LifeFlow. Finally we would like to thank Ben Shneiderman for his contribution to the design of LifeFlow, and Nikola Ivanov and Michael VanDaniker for their insightful comments and support.

REFERENCES

1. Motor Vehicle Traffic Crashes As a Leading Cause of Death in the United States, 2005. Traffic Safety Facts. Research Note. National Highway Traffic Safety Administration, <http://www-nrd.nhtsa.dot.gov/Pubs/810936.PDF>. April 2008. Accessed on July 16, 2009.
2. Traffic Safety Facts 2006. A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System. National Highway Traffic Safety Administration
3. Pack, M. Visualization in Transportation: Challenges and Opportunities for Everyone, *IEEE Computer Graphics and Applications*, pp. 90-96, July/August, 2010
4. Wongsuphasawat, K., T. Wang, C. Plaisant and B. Shneiderman. Finding Patterns in Temporal Data, *HCIL Open House 2010*, www.cs.umd.edu/hcil/temporalviz
5. Tufte, E., *The Visual Display of Quantitative Information*, 2nd Edition, (2001) Graphics Press
6. Fredrikson, A., C., North, C., Plaisant and B., Shneiderman. Temporal, Geographical and Categorical Aggregations Viewed through Coordinated Displays: A Case Study with Highway Incident Data, *Proceeding of the Workshop on New Paradigms in Information Visualization and Manipulation*, Kansas City, Missouri, November 6, 1999 (in conjunction with ACM CIKM'99), ACM, 26-34.
7. Wongsuphasawat, K., D., Filippova, M., VanDaniker, M., Pack and A. Olea. Visual Analytics for Transportation Incident Datasets in *Proc. Transportation Research Board 88th Annual Meeting*, Washington, DC, USA, 11-15 January 2009.
8. Pack, M. L., K., Wongsuphasawat, M., VanDaniker and D. Filippova. ICE--visual analytics for transportation incident datasets. In *Proceedings of the 10th IEEE international Conference on information Reuse & integration* (2009) 200-205.
9. Plaisant, C., P. Tarnoff, S., Keswani, A., Saraf and A., Rose. Understanding Transportation Management Systems Performance with a Simulation-Based Learning Environment, *Proceeding of Conference on Intelligent Transportation Systems' 99*,

ITS'99, Washington, DC, ITS America, Washington DC, <http://www.itsa.org> (CD ROM proceedings) 1999.

10. Plaisant, C., R., Mushlin, A. Snyder, J., Li, D., Heller and B., Shneiderman. LifeLines: using visualization to enhance navigation and analysis of patient records. *Proceedings American Medical Informatics Association Symposium*. 1998: 76–80.
11. Ellison L. J., N., Ivanov. Real-Time and Historic Incident Visualization Using Timelines *86th Annual Meeting of the Transportation Research Board Compendium of Papers CD-ROM* January 21st-25th, Washington DC, 2007
12. Wang, T. D., C., Plaisant, A.J., Quinn, R., Stanchak, S., Murphy, and B., Shneiderman. 2008. Aligning temporal data by sentinel events: discovering patterns in electronic health records. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 457-466
13. Wongsuphasawat, K. and B., Shneiderman. Finding Comparable Temporal Categorical Records: A Similarity Measure with an Interactive Visualization, *Proceedings of IEEE Visual Analytics*, 2009.
14. Fekete, J.D., "The InfoVis Toolkit," Information Visualization, IEEE Symposium on, pp. 167-174, 2004 *IEEE Symposium on Information Visualization* (InfoVis 2004), 2004
15. Bederson, B., B. Shneiderman and M., Wattenberg. Ordered and Quantum Treemaps: Making Effective Use of 2D Space to Display Hierarchies, *ACM Transactions on Graphics*, 21, 4, 2002, 833-854.