Algorithms for Identifying and Ranking Bottlenecks Using Probe Data

Andrew Lund
Senior DBA, CATT Lab
301.405.2874
Alund1@umd.edu

Michael L. Pack
Director, CATT Lab
301.405.0722
PackML@umd.edu

Catherine Plaisant
Faculty/UX Lead/HCIL Member
301.405.2768
cplaisant@cs.umd.edu

Mark L. Franz, Ph.D.
Faculty Specialist, CATT Lab
301.314.0422
mfranz1@umd.edu

University of Maryland
CATT Laboratory
5000 College Ave. Suite 2206
College Park, MD 20742

Corresponding Author: Andrew Lund

Word/Figure Count: 5,230 words + 6 Figures and Tables = 6,730 total words

Revised Submission Date: November 11, 2016

ABSTRACT

Transportation officials are constantly under scrutiny to improve congestion on the nation’s roadways. Identifying where congestion exists is getting easier with the advancement of probe-based speed and travel time data from companies like HERE, INRIX, and TomTom; however, identifying the source of congestion and the actual bottleneck locations, can be significantly more challenging. In this paper, we discuss the dynamic nature of congestion and bottlenecks—how they can grow and shrink over time; how two or more bottlenecks can merge with one another; how a single bottleneck can split and become two separate bottleneck events; and how bottlenecks can be both recurring and non-recurring. We propose an algorithm for identifying bottleneck locations along with terminology for keeping track of the various components that make up a bottleneck.

INTRODUCTION/BACKGROUND

Transportation officials are constantly under scrutiny to improve congestion on the nation’s roadways. Identifying where congestion exists is getting easier with the advancement of probe-based speed and travel time data from companies like HERE, INRIX, and TomTom. These data providers typically report speeds on nearly every segment of a roadway at 1-minute intervals. These data providers are also working to increase the spatial resolution of their data—fine-tuning the segment lengths down to just a few hundred feet in some cases. As both temporal and spatial resolution of these data sources gets better and better, it is becoming much easier to have a ubiquitous picture of where congestion on the roadway is happening. However, identifying the head of the bottleneck location can be significantly more challenging. Note that throughout this paper, the term “bottleneck” refers to the congestion that plagues a section of the road and the term “head of the bottleneck” refers to the location on the road where the congestion originates.

Bottlenecks are dynamic. They grow, shrink, and can merge with neighboring bottlenecks over their lifetime. Identifying the head of a bottleneck is critical to identifying problem areas on the roadway. Here we discuss the dynamic nature of congestion and bottlenecks—how they can grow and shrink over time; how two or more bottlenecks can merge with one another; how a single bottleneck can split and become two separate bottleneck events; and how bottlenecks can be both recurring and non-recurring. We propose an algorithm for identifying bottleneck locations along with terminology for keeping track of the various components that make up a bottleneck.

Dynamic Bottlenecks

Existing bottleneck identification methods are often quite good at identifying the general locations of the worst bottlenecks within a road network, and when attempting to rank bottleneck locations by aggregating the attributes of the bottlenecks at each location, known problem areas will usually show up at the top of the list. This is a good starting point, but due to the dynamic nature of congestion, especially along long stretches of heavily trafficked roadways, existing bottleneck identification methods have a number of challenges and limitations when it comes to analyzing and ranking bottlenecks.

Figure 1 shows several congestion graphs depicting congestion patterns in which existing bottleneck identification methods may not accurately reflect conditions. Each graph displays traffic conditions at locations *A* - *D* during sequential time periods *t1* – *t4*. A red cell indicates congested conditions, while a green cell indicates non-congested conditions. Traffic travels along the road from top to bottom, so a traveler will encounter location *D* first and proceed to *A*.

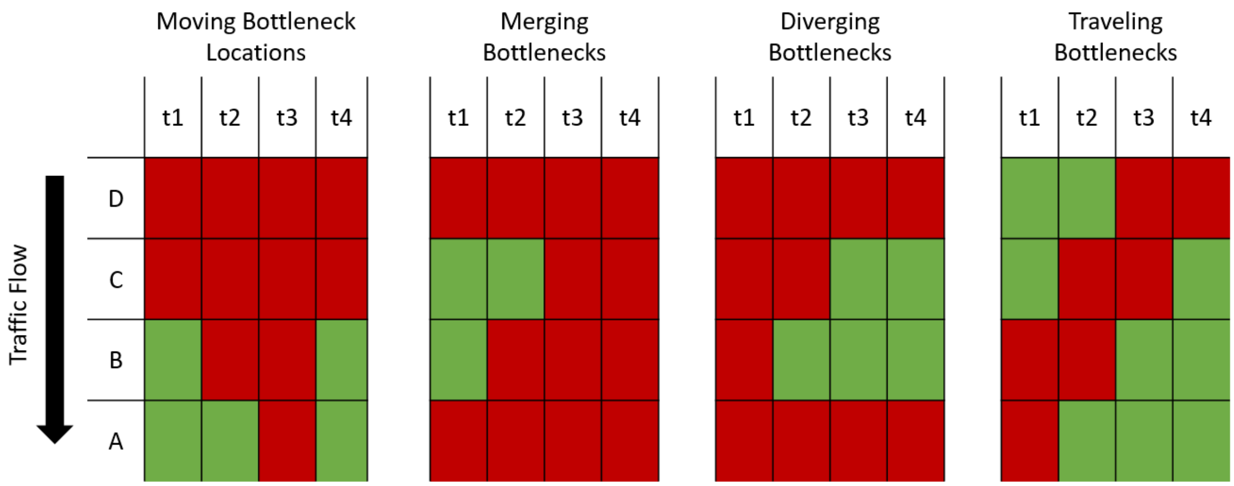

FIGURE 1 Examples of challenging congestion patterns for bottleneck identification.

Moving Bottleneck Locations

Over the life of a bottleneck, the location of congestion on the roadway may change by moving either downstream or upstream of the original location. These changes in location may be relatively short-lived compared to the overall duration of the bottleneck and therefore arguably less significant with respect to identifying the actual problem location causing the congestion. In cases where the congestion moves downstream of the original location temporarily, the bottleneck should then be defined by the new furthest downstream location even if congestion returns to its original location.

Merging Bottlenecks

In many cases, two or more bottlenecks may merge if congestion conditions grow to include any road segments that previously separated the bottlenecks. Using existing bottleneck identification methods, the final bottleneck would be defined at the furthest downstream location among all of the merged bottlenecks. Depending on the characteristics of the merging bottlenecks, the reported bottleneck location may be misleading if the upstream bottleneck(s) that were merged had a larger impact on the roadway than the furthest downstream bottleneck. This limitation is similar to the "Moving bottleneck locations" issue described above, but may result in more severe misrepresentations of the data.

Diverging Bottlenecks

In addition to congested conditions growing and causing multiple bottlenecks to merge over time, congestion may also clear somewhere in the middle of an existing bottleneck. In this case, there will now be two groups of consecutive congested road segments that should be tracked as separate bottlenecks. However, existing bottleneck identification methods will preserve the link between the two bottlenecks and will define both as being located at the same location using the furthest downstream location of the two. Similar to the "Moving bottleneck locations" and "Merging bottlenecks" issues described above, this can lead to severe misrepresentations of the data depending on how the congestion behaves after the original bottleneck diverges.

Travelling Bottlenecks

Slow-moving vehicles, such as line painting trucks or vehicles moving extremely oversized loads, may sometimes cause a “travelling bottleneck” condition in which a short stretch of congestion follows the slow vehicle for a considerable distance along a roadway (*1*). Using existing bottleneck identification methods, situations like this may be identified as a single bottleneck that affected the entire roadway on which the slow vehicle was traveling, even if no single location along that roadway is congested for more than a few minutes at a time.

RELATED WORK

The identification of bottlenecks has been an important topic of study since congestion has been observed on roads. Early “floating cars” studies were eventually replaced by algorithms using point data from detector loops embedded in the pavement (*2*,*3*), often harnessing the help of simulations (*4*). Recent techniques use continuous probe data collected from Global Positioning Systems installed in vehicle fleets (*5*,*6*).

Even detecting whether a road segment is congested or not is not trivial. For example, Elhenawy et al. showed the benefit of combining speed distributions in congestion and free-flow (*7*). In dense road networks the identification of trajectories (*8*) and bottlenecks covering multiple roads can also be challenging (*9*).

A large amount of work has focused on real time management issues, e.g. Jin et al. use pattern matching in traffic data to detect incidents in real time (*10*), others study bottlenecks to provide travel time prediction (*11*), or to inform variable speed advisories (*11*).

Identifying bottlenecks within a road network is typically accomplished by scanning any speed, travel time, or traffic flow data available on each road and looking for groups of consecutive road segments that are all exhibiting congested conditions at the same time. The specific criteria required for a road segment to be considered congested may be defined differently based on the goals of the analysis or the characteristics of the data being used (*12,13*). A common approach for defining the congested criteria is to set a specific threshold value for observed travel speeds in relation to a reference travel speed, usually either a posted speed limit or a calculated free-flow speed.

After identifying a group of consecutive congested road segments, the conditions at this location may be tracked over time to determine if it is actually a bottleneck. An example of the observed life of a bottleneck is illustrated in Figure 2.

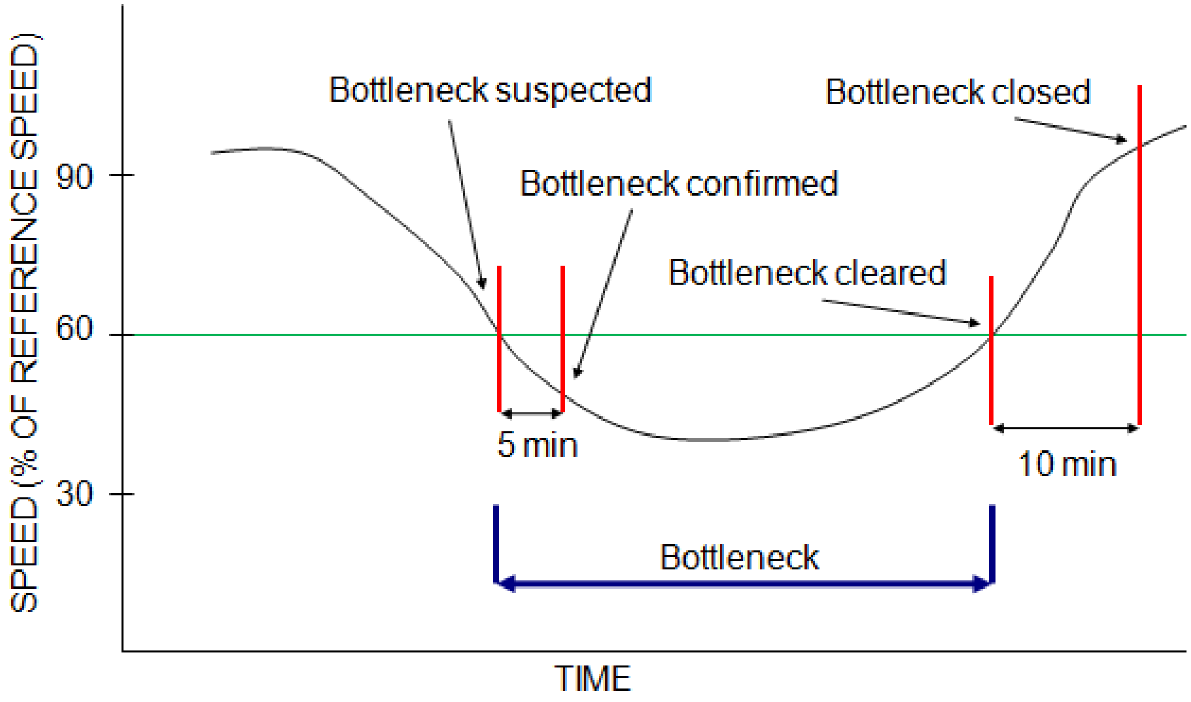

FIGURE 2 Example of observed life of a bottleneck.

Figure 2 uses an observed speed equal to 60% of the reference speed as the criteria for identifying congested road segments. Once congestion has been identified, it then requires that conditions remain in a congested state for 5 minutes before the location is confirmed as a bottleneck. The bottleneck remains active until the congested conditions clear and remain cleared for 10 minutes. These 5- and 10-minute buffer times at the beginning and end of the bottleneck are meant to address false positives within the data being used for analysis.

Following the identification of bottlenecks on a road network, each bottleneck should have a location defined, typically using the furthest downstream congested location (*12*). Any additional attributes of the bottleneck available in the data set being used may be assigned to each bottleneck to assist in comparing or ranking bottlenecks, such as duration, length of the roadway affected, or number of vehicles affected.

THE NEW ALGORITHM

The challenges with existing bottleneck identification methods are primarily rooted in how bottleneck locations are defined by using the furthest downstream location regardless of how congestion evolves. To address these challenges, we introduce a new congestion tracking method, new terminology, and new visualizations for exploring the data. The new congestion tracking method gives more focus to individual locations causing congestion, allowing for more accurate bottleneck location identification and ranking. For this algorithm, we introduce several new terms:

- Occurrence – Congestion, whose head is at a given point on the road at a single point in time

- Element—Congestion, whose head is at a given point on the road, that can change in length over time
- Blob—a collection of spatially and temporally adjacent congestion elements

Occurrences

The foundation of the proposed bottleneck identification method is built on the same basic concepts as existing methods: a traffic data set is analyzed at each reading interval to identify groups of consecutive congested road segments, or occurrences. An occurrence is formally defined as an area of congestion whose head is at a given point on the road at a single point in time. As with existing bottleneck identification methods, the head of congestion is defined as the furthest downstream location.

Upon identification, each occurrence is assigned a set of attributes derived from the data set being analyzed. These attributes include information such as head location, the date and time at which the occurrence was observed, information about the roadway segments that were congested as part of the occurrence (such as a set of road segment identifiers), and an impact value. The impact of an occurrence may be calculated in whatever way is deemed appropriate for the analysis being done as long as the calculation of the impact value is identical for every occurrence present. One simple suggestion for calculating impact is to sum the lengths of every road segment included in the occurrence. If desired, these basic impact values may be weighted to better reflect the severity of the congestion by using vehicle volumes or observed speed deviations from a reference speed. In some cases, assigning multiple impact values using different calculations for each occurrence may be useful.

Elements

Once a single reading interval in the traffic data set is analyzed to identify all of the occurrences, an identical analysis is performed on the next sequential reading available. Occurrences from sequential data readings that share identical head locations are combined into elements. However, occurrences that do not share an identical head location to any occurrence from the previous reading interval will become a new element. An element is formally defined as an area of congestion whose head is at a given point on the road that can change in length and/or severity over time.

Elements are assigned a set of attributes that are derived from the occurrences they are comprised of, including a head location, a start date and time (derived from the earliest occurrence), information about the roadway segments that were congested as part of the element (the unioned set of congested roadway segments for all occurrences), and an impact value (the sum of the individual occurrence impact values).

Each sequential reading interval analyzed will produce its own set of occurrences that will be used to either update the attributes of existing elements from the previous reading interval that share identical head locations or create new elements. Elements that exist in one reading interval but do not have an occurrence in the next sequential reading with an identical head location are assigned end date and time attributes.

The final state of an element's attributes provides a general indication of the total effect congestion at that location had on the roadway. Because of this, elements will serve as the base unit for identifying and ranking bottleneck locations under the new method. This will be discussed in more detail in the "Bottleneck Ranking" section below.

Blobs

A key difference between existing bottleneck identification methods and the new method described in this paper is the focus on individual locations at which congestion is located. A bottleneck defined at one location by an existing method may include congestion caused by conditions at one or more additional locations upstream. While this approach often does not produce the results you may expect (see section Challenges and limitations of existing methods above), there are benefits to preserving the relationship between the different congested locations. To address this, related elements are grouped into a blob. A blob is formally defined as a collection of spatially and temporally adjacent congestion elements.

Similar to how elements are built from a collection of occurrences, blobs are built from a collection of elements. The main difference is that while all occurrences that make up an element must share an identical head location, a blob includes all elements that move to, merge into, or diverge from each other.

Blobs are assigned attributes that are derived from the elements they are comprised of, including a head location (the furthest downstream head location from all elements), a tail location (the furthest upstream head location from all elements), a start date and time (the earliest start time of all elements), an end date and time (the latest end time of all elements), information about the roadway segments that were congested as part of the blob (the unioned set of congested roadway segments for all elements), and an impact value (the sum of the individual element impact values). This information may be useful to help identify stretches of road that have multiple locations contributing to congestion that regularly impact each other, or to filter out short-lived congestion that is either insignificant or caused by a momentary lapse in data quality.

BOTTLENECK RANKING

Under the new method of tracking congestion described above, individual locations along a roadway network may be ranked against each other by summing up the impact values for all elements occurring at each location and then ordering the locations using these summations. Locations with the highest total impact value are determined to be the worst bottleneck locations and will represent areas of recurring congestion or areas that experienced severe non-recurring events during the time period being analyzed. Because elements can never change location, the total impact value calculated from all the elements at a given location provides a more accurate depiction of what impact that specific location had on the roadway. This is a significant improvement in the precision of bottleneck analysis over existing methods as it removes the problem of congestion caused by locations further upstream being attributed to the furthest downstream location.

One noted side effect of the increased location precision in bottleneck identification under the new method is that stretches of road with multiple locations close together that all contribute to recurring congestion on their own may end up being ranked notably lower in a wider-area bottleneck ranking than they would be under existing methods. The reason for this is that these areas would usually all be grouped together under a single ranked bottleneck location with each of their impacts combined into the total, while the new method will treat each location separately and therefore potentially splitting the impacts significantly depending on how many contributing locations are included. If desired, additional rules may be defined at the time of element aggregation to group nearby locations into a single ranking. If this is done, all of the impact values for elements located at any of the specific locations in the group would simply be aggregated together as if they were all located at the same location.

EXAMPLE EXERCISE IN IDENTIFYING AND RANKING BOTTLENECKS

To help demonstrate the proposed bottleneck identification and ranking method, the following exercise will walk through a simple example using a grid representation of a contour-line style traffic congestion graph (for an example of this style of graphic, see Lund and Pack (*5*)).

	t1	t2	t3	t4	t5	t6	t7	t8
H	🟥	🟥	🟥	🟥	🟥	🟥	🟥	🟥
G	🟥	🟥	🟥	🟥	🟥	🟥	🟥	🟥
F	🟩	🟩	🟥	🟥	🟩	🟩	🟩	🟩
E	🟩	🟥	🟥	🟩	🟩	🟩	🟩	🟩
D	🟥	🟥	🟥	🟥	🟩	🟩	🟩	🟩
C	🟥	🟥	🟥	🟥	🟩	🟩	🟥	🟩
B	🟩	🟥	🟥	🟩	🟩	🟩	🟥	🟥
A	🟩	🟩	🟩	🟩	🟩	🟩	🟥	🟥

FIGURE 3 Traffic congestion graph.

In Figure 3, each column represents a point in time at which a vehicle probe speed reading is measured (identified as t1 through t8), and each row represents a road segment along a continuous stretch of road (identified as A through H). Traffic flows downward, so a motorist traveling this stretch of road would encounter segment H first and proceed to travel towards segment A. A red cell indicates congested conditions, while a green cell indicates non-congested conditions. To simplify calculations, each road segment A - H is assumed to have a length of 1 unit. Each timestamp t1 – t8 will be analyzed in order, tracking occurrences, elements, and blobs as described earlier in this document.

t1 – During the initial reading interval t1, there are two groups of consecutive congested road segments, or occurrences, located at C and G. Both occurrence span two road segments (C-D and G-H, respectively), so they will each have an impact of 2. Since there are no existing elements at this time, there will be two new elements created consisting of the identified occurrences. There are also no existing blobs at this time, so the identified elements begin as part of separate blobs. The blob at C will be identified as b1 and the blob at G will be identified as b2.

t2 – At reading t2 an occurrence no longer exists at C as congestion has shifted downstream to B. A new occurrence is identified at B with an impact of 4, and a new element is created from that occurrence since there were no elements located at B during the previous reading. Because the new element at B is spatially and temporally adjacent to the element at C from t1, the new element will belong to the same blob b1. The existing element at C is now closed. An occurrence at G is identified again, so the attributes of the existing element from the previous reading will be updated by adding the new occurrence's impact value of 2, bringing its total impact to 4.

t3 – Congestion located at B has grown further upstream and has merged with the previous congestion at G. The existing element at B is updated by adding the new occurrence's impact value of 7, bringing its total impact to 11. The existing element at G is now closed. Since the new occurrence at B overlaps two previous occurrences belonging to separate blobs (b1 and b2), the two blobs are now merged together into a single blob b1.

t4 – There are now two separate occurrences again, one at C and one at F. The existing element at B is now closed, and a new element is created at C with an impact value of 2. Another new element is also created at F with an impact value of 3. Both new elements belong to blob b1 since they are spatially and temporally adjacent to the previous element at B. Note that there are now two active elements belonging to the same blob.

t5 – Congestion at C has cleared, so the existing element at C is closed. The congestion at F has now moved upstream, so the existing element at F is closed and a new occurrence and element is created at G with an impact value of 2 and belonging to blob b1.

t6 –The congestion at G remains, so the attributes of the existing element at G are updated by adding the new occurrence's impact value of 2, bringing its total impact to 4.

t7 – An occurrence is identified at A that is not spatially adjacent to any elements from the previous reading, so a new element will be created with an impact value of 3 and belonging to a new blob b3. The congestion at G remains unchanged, so the attributes of the existing element are updated by adding the new occurrence's impact value of 2, bringing its total impact to 6.

t8 – Congestion at both A and G remains and both occurrences have an impact value of 2, so the corresponding elements are updated by increasing the total impact to 5 and 8, respectively.

Final Element Attributes

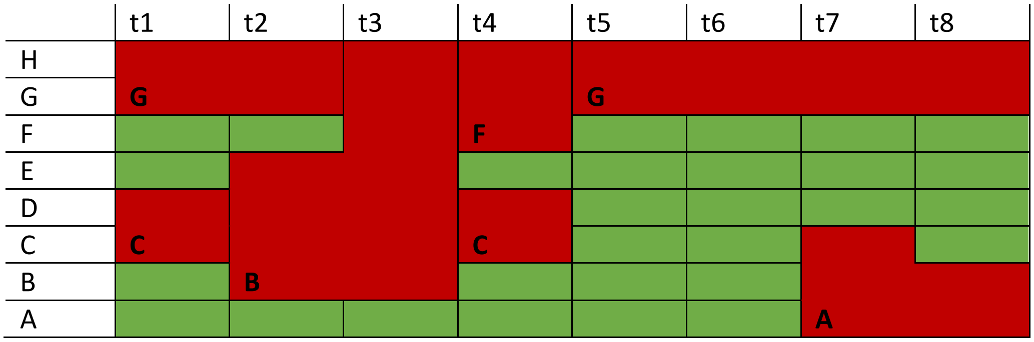

Location	Start	End	Total Impact	Segments Affected	Blob
C	t1	t2	2	CD	b1
G	t1	t3	~~2~~ 4	GH	~~b2~~ b1
B	t2	t4	~~4~~ 11	BCDEFGH	b1
C	t4	t5	2	CD	b1
F	t4	t5	3	FGH	b1
G	t5	t9	~~2~~ ~~4~~ ~~6~~ 8	GH	b1
A	t7	t9	~~3~~ 5	ABC	b3

FIGURE 4 Example exercise elements and element attributes.

Figure 4 shows a similar contour-line style congestion graph as Figure 3, only this time the elements identified while analyzing the data using the proposed bottleneck identification method have been outlined and labeled with their location's name in the lower left corner. Underneath the congestion graph, the final attributes for each element are shown in a table format. Attribute values that were assigned to an element at one reading but were updated in subsequent readings are shown with a strikethrough style for illustrative purposes. For example, the element at location G starting at t1 originally had an impact value of 2 which was later updated to 4 during analysis of t2. This element also began as part of its own blob b2, but was later merged into blob b1 during analysis of t3.

Bottleneck Ranking Example Exercise

After identifying all of the elements and associated blobs within a traffic data set, the resulting information may be used to identify bottleneck locations along the roadway network and rank them according to their overall impact on traffic. To do this, all elements with identical locations are aggregated together, summing up their total impacts. Table 1 (top) shows the bottleneck locations identified in the example exercise in decreasing order of total impact.

TABLE 1 Bottleneck Rankings for Example Exercise Using Element Aggregates (top) and Example Exercise Blob Attributes (bottom)

Bottleneck Location	Total impact
G	12
B	11
A	5
C	4
F	3

Blob	Total impact	Segments Affected	Start Time	End Time
b1	30	BCDEFGH	t1	t11
b3	5	ABC	t9	t11

Upon a cursory visual analysis of the congestion graph, these results may not be immediately obvious as congestion occurring at segment B appears to have had the worst congestion based on the longer length of congested segments occurring there; however, a more detailed analysis reveals that segment G actually had a slightly greater impact on traffic due to its consistently congested conditions over a longer duration. Existing bottleneck tracking methods would have attributed the entire impact of blob b1 to location B since it is the furthest downstream location, and the significant impact of location G would be lost.

Using element aggregation to identify and rank individual bottleneck is an effective way to pinpoint locations that have the largest impact on a roadway network, but on occasion it may still be desirable to analyze congestion in a way that is more similar to existing bottleneck identification methods. For these cases, looking at blob attributes will be useful.

Table 1 (bottom) shows the various attributes assigned to the two blobs identified in the example exercise above. By comparing the impacts of different blobs, particularly severe instances of congestion can be identified which may suggest a major traffic incident or other traffic event. While there are only two blobs in this sample exercise, an analysis over a larger road network may yield only a handful of abnormally high-impact blobs out of hundreds of smaller ones.

REAL WORLD APPLICATION

To illustrate the advantages of the new bottleneck identification and ranking algorithm over existing methods, a real world application is presented using field data along the northbound segment of I-95 in Maryland approaching the MD-100 interchange. This site was selected because it regularly ranks as one of the top bottlenecks in the state of Maryland (*14*) and is known to have multiple nearby locations that regularly contribute to congestion. Using a GPS probe data set for the month of October 2016, this location was evaluated with implementations of both existing algorithms and the new algorithm proposed in this document. Table 2 summarizes the application of both bottleneck algorithms on the study segment.

TABLE 2 Existing and New Algorithm Application Results

	Location	Impact	Average max length (miles)	Average duration	Total Duration
Existing Algorithm	I-95 N @ MD-100/EXIT 43	69,742	7.69	2 h 01 m	6 d 7 h 15 m
	I-95 N @ MD-175/EXIT 41	4,378	3.9	1 h 06 m	0 d 18 h 42 m
	I-95 N @ MD-32/EXIT 38	1,044	2.96	32 m	0 d 5 h 52 m
New Algorithm	I-95 N @ MD-100/EXIT 43	15,904	4.83	1 h 43 m	2 d 05 h 40 m
	I-95 N @ MD-175/EXIT 41	10,396	3.76	1 h 43 m	2 d 05 h 15 m
	I-95 N @ MD-32/EXIT 38	3,672	2.74	36 m	0 h 18 h 41 m

When comparing the results in Table 2, it is important to note that the impact values are calculated differently between the existing algorithm and new algorithm implementations and therefore are not directly comparable. However, the relative size of the impact between locations within an implementation serves an identical purpose of comparing the severity of the bottleneck locations. With this in mind, it is clear that the existing algorithm implementation attributes a significant majority of the congestion to the I-95 @ MD-100 bottleneck while potentially underestimating the impact of upstream locations. In fact, the impact at the I-95 @ MD-100 bottleneck was 16 times larger than that of the I-95 @ MD-175 location and 67 times larger than that of the I-95 @ MD-32 location. Here, the length of the affected roadway segment as well as the duration of congestion do not reflect the observed congestion patterns at these segments of I-95. Recall that the existing algorithm identifies bottlenecks as sustained speed drops and assigns the associated impacts to the furthest downstream segment without considering the complex interaction of congestion resulting from multiple bottlenecks. In doing so, the existing algorithm may underestimate the impact of upstream bottlenecks on a heavily congested corridor due to congestion at these locations merging with the congestion of downstream bottlenecks.

To overcome this potential shortcoming, the new bottleneck algorithm evaluates the spatial and temporal evolution of congestion resulting from the formation of multiple bottlenecks along a congested corridor. In evaluating the new algorithm at the I-95 analysis site, the congestion is more realistically represented across the three bottleneck locations. Here, the impact factor at the I-95 @ MD-100 bottleneck is only 1.5 times larger than that of the I-95 @ MD-175 location and 4.3 times larger than that of the I-95 @ MD-32 location. In a similar fashion, the length of the affected roadway and total duration better reflect observed congestion patterns and the contribution of each bottleneck location to the overall congestion on this portion of I-95. Note that if the congestion was in fact primarily caused by the I-95 @ MD-100 bottleneck the new algorithm would have produced results similar to those from the existing algorithm.

CONCLUSIONS AND FUTURE WORK

The proposed algorithm is tailored towards probe-based data. The addition of ubiquitous real-time volume if and when it becomes available would make it easier to incorporate user delay costs as a component of the bottleneck ranking methodology. There are still a number of edge-cases that this algorithm may not be able to capture fully yet—especially those related to extremely short road segments that fluctuate between a congested and non-congested state due to traffic signals or poor data quality. In addition, the development of novel online bottleneck visualization tools is underway. These visualizations will allow transportation analysts to better understand and communicate their bottleneck analysis findings.

Our proposed bottleneck identification methodology has shown significant improvements over existing methods—specifically with respect to reducing errors caused by merging and separating bottlenecks, as well as more accurately identifying locations contributing to congestion. These improvements have the potential to assist transportation analysts in understanding the complex nature of congestion and better prioritize congestion mitigation projects.

REFERENCES

1. Munoz, J. C., and C. F. Daganzo. Moving Bottlenecks: A Theory Grounded on Experimental Observation. In 15th International Symposium on Transportation and Traffic Theory (M. A. P. Taylor, ed.), Pergamon-Elsevier, Oxford, United Kingdom, 2002, pp. 441–462.
2. Chen, C., A. Skabardonis, and P. Varaiya. Systematic Identification of Freeway Bottlenecks. In Transportation Research Record: Journal of the Transportation Research Board, No. 1867, Transportation Research Board of the National Academies, Washington, D.C., 2004, pp. 46–52
3. Bertini, R. L., and A. M. Myton. Use of Performance Measurement System Data to Diagnose Freeway Bottleneck Locations Empirically in Orange County, California. In Transportation Research Record: Journal of the Transportation Research Board, No. 1925, Transportation Research Board of the National Academies, Washington, D.C., 2005, pp. 48–57
4. Ban, X., Chu, L., & Benouar, H., Bottleneck identification and calibration for corridor management planning. *Transportation Research Record: Journal of the Transportation Research Board*, 1999 (2007) 40-53
5. Lund, A., & Pack, M. (2010). Dynamic wide-area congestion and incident monitoring using probe data. *Transportation Research Record: Journal of the Transportation Research Board*, (2174), 1-9.
6. Zhao, W., McCormack, E., Dailey, D. J., & Scharnhorst, E., Using truck probe GPS data to identify and rank roadway bottlenecks. *Journal of Transportation Engineering*, *139*,1 (2012) 1-7.
7. Elhenawy, M., & Rakha, H. A., Automatic Congestion Identification with Two-Component Mixture Models. *Transportation Research Record: Journal of the Transportation Research Board*, 2489 (2015) 11-19
8. von Landesberger, T., Brodkorb, F., Roskosch, P., Andrienko, N., Andrienko, G., & Kerren, A. (2016). MobilityGraphs: Visual Analysis of Mass Mobility Dynamics via Spatio-Temporal Graphs and Clustering. *Visualization and Computer Graphics, IEEE Transactions on*, *22*, 1 (2015) 11-20.
9. Wang, X., Lu, G., Chen, P., & Wang, Y., Bottleneck Identification on Beijing Road Network Based on Reliability. In *15th COTA International Conference of Transportation Professionals* (2015).
10. Jin, Y., Dai, J., & Lu, C. T., Spatial-temporal data mining in traffic incident detection. In *Proc. SIAM DM 2006 Workshop on Spatial Data Mining*, 5 (2006).
11. Lu, X. Y., Shladover, S. E., Jawad, I., Jagannathan, R., & Phillips, T., Novel Algorithm for Variable Speed Limits and Advisories for a Freeway Corridor with Multiple Bottlenecks. *Transportation Research Record: Journal of the Transportation Research Board*, 2489 (2015) 86-96.
12. Chen, C., A. Skabardonis, and P. Varaiya. Systematic Identification of Freeway Bottlenecks. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1867,* Transportation Research Board of the National Academies, Washington, D.C., 2004, pp. 46–52.
13. Bertini, R., & Myton, A,. Use of performance measurement system data to diagnose freeway bottleneck locations empirically in Orange County, California. *Transportation Research Record: Journal of the Transportation Research Board*, 1925 (2005) 48-57
14. Mahapatra, S., Wolniak, M., Beckett, E., and Sadabadi, K.V. 2015 Maryland State Highway Mobility Report. Maryland Department of Transportation, Hanover, MD. December 2015.