

Report no. VNTSC-TFM-08-11

A New Method for Determining a Sector Alert

Eugene Gilbo Rick Oiesen

29 September 2008

Prepared for Federal Aviation Administration Office of System Operations Programs 800 Independence Ave., SW Washington, DC 20591

Prepared by Volpe National Transportation Systems Center Traffic Flow Management Division 55 Broadway Cambridge, MA 02142

Acknowledgments

Within the FAA, we would like to thank Midori Tanino and Ved Sud, who have supported and encouraged this work. Traffic managers who have critically read and commented on this proposal include Archer Davis of Boston Center, Mark Evans of Cleveland Center, Philip Bassett of Jacksonville Center, Curt Kaler of Minneapolis Center, Leo Prusak of the LaGuardia Tower, and William Roth of Boston Center, This report reflects the discussions and feedback that they have provided, but they should not be blamed for anything in this report.

Within Volpe, we would like to thank Mary Costello for editing the report.

Executive Summary

The Traffic Flow Management System (TFMS) currently declares an alert for any 15-minute interval in which the predicted demand exceeds the Monitor/Alert Parameter (MAP) for any airport, sector, or fix. For a sector, TFMS predicts the demand for each minute, and TFMS uses the demand of the peak minute in a 15-minute interval to decide whether to declare an alert for the entire 15-minute interval. Using the peak demand from a single minute to declare alerts has been criticized for three reasons. First, the demand for a single minute is a flawed measure of workload for the entire 15-minute interval. Second, using demand for a single minute leads to instability; that is, slight fluctuations in demand from minute to minute can lead to alerts flickering on and off. Third, the interval that is alerted depends on the arbitrary, 15-minute boundaries.

To deal with these problems with the current method of declaring alerts, and to develop a method of declaring alerts that matches more closely the intuition of traffic managers about what patterns of demand represent a potential problem, this report proposes using information about both the *magnitude* and *duration* of excess demand to determine if there is an alert. To make this determination, this report defines three parameters: **a**, **b**, and **A**. By choosing particular values for these parameters, TFMS could control the magnitude and duration of excess demand that is required to trigger an alert. The thinking expressed by traffic managers is that a single minute of slightly excess demand should not be enough to trigger an alert. Therefore, a parameter **a** is defined as the number of minutes of excess demand that must occur for there to be an alert. (Excess demand means that the demand for a minute exceeds the MAP.) For example, if **a** = 3, then at least three minutes of excess demand are needed for an alert. The next question is how bunched in time these three minutes need to be. The proposal is that at least **b** consecutive minutes of normal, i.e., not excess, demand, are enough to prevent or end an alert. The larger **a** is set, and the smaller **b** is set, the harder it is for an alert to occur.

Traffic managers also state that while one minute of slightly excess demand is not enough to justify an alert, even one minute of significant excess demand is enough. Therefore, the parameter Δ is used to determine how large the demand for a single minute should be to trigger an alert. If demand for a minute is greater than MAP + Δ , an alert is declared. This is called a short-term alert since it can arise from the demand for a single minute, while the type of alert explained in the previous paragraph is called a long-term alert.

Discussions with traffic managers have led to the conclusion that this proposed new method of determining alerts holds promise, but it would be premature to say that it is ready to be implemented in TFMS. Therefore, it is proposed that a prototype of this method be implemented so that traffic managers can try it and also so that analysis can be carried out. This prototype could be used, for example, to try different values for the various parameters, and it could then be seen how this affected the frequency and stability of alerts. A prototype would also allow traffic managers to evaluate the usefulness of the concepts of short- and long-term alerts and to determine whether any alternate definition might be more useful. Some questions are whether the parameters should have the same value for the country as a whole, whether they should have the same values for every sector in each center, or whether they should be separately determined for every sector. Finally, there are user interface questions, such as how these two different types of alerts should be displayed.

In summary, the goal is to allow TFMS to reliably identify patterns of demand that signal a potential problem. An alert would be used to call this potential problem to the attention of a traffic manager, who could then look at the situation in detail and decide if any action is needed.

Table of Contents

1.	Introduction and Purpose of this Report			1
2.	Problems with the Current Definition of Sector Alerts			1
3.	Logic Underlying the Proposed Approach			2
4.	Propo	Proposed New Approach for Declaring Sector Alerts		3
	4.1 Loi		g-term Alert	3
	4.2	Short	term Alert	4
5.	Summ	ary ai	nd Next Steps	7
Appendix A		A	Duration Limits for Short- and Long-term Alerts	9
Ap	pendix	B	An Option for Displaying Short-term Alerts: Showing Insignificant Overload	11
Appendix		C '	The Unified Sector Alert Algorithm	12

1. Introduction and Purpose of this Report

The Traffic Flow Management System (TFMS) predicts the volume of air traffic for airports, sectors, and fixes to identify potential congestion problems. TFMS's Monitor/Alert function compares predicted traffic demand to the Monitor/Alert Parameter (MAP) and issues alerts wherever and whenever demand is predicted to exceed the MAP. This helps TFM specialists decide when action is needed to deal with a potential congestion problem. (Note: In June 2008, TFMS replaced the Enhanced Traffic Management System (ETMS). TFMS, however, uses exactly the same methods for determining alerts as ETMS.)

Although the general alerting rule is the same for airports, fixes and en route sectors and is based on comparison of predicted demand with MAP, TFMS defines traffic demand for sectors differently than for airports and fixes. Traffic demand for airports and fixes is measured by the aggregate number of flights per 15-minute interval, i.e., number of arrivals or departures at an airport per 15-minute interval or number of flights crossing a fix per 15-minute interval. These aggregate 15-minute demand counts are used in determining airport or fix alert status. For sectors, TFMS uses a different approach to determining sector demand and sector alert status.

Currently, TFMS uses the following algorithm to determine if a sector is alerted in a 15-minute interval.

- For a 15-mintes interval, TFMS predicts the demand for each minute in that interval, i.e., the number of flights that will be in that sector during that minute.
- TFMS compares the demand for each minute with the MAP for that sector.
- If the demand for any one minute exceeds the MAP, then the entire 15-minute interval containing that minute is alerted.

Stated another way, the maximum demand for any minute during a 15-minute interval is taken to be the demand for that interval; it is clear, however, that the peak aircraft count during a 15-minute interval may not adequately reflect the true demand on that sector during the 15-minute interval.

Thus, Monitor/Alert treats sectors differently from airports and fixes. Whereas the aggregate demand for the entire 15-minute is taken to be the airport and fix demand for that interval, for the sector, the demand for the peak minute of a 15-minute interval is currently taken to be the sector demand for that 15-minute interval, and this is the figure that is compared to the MAP and determines an alert status.

The purpose of this report is to propose a new method of defining sector alerts.

2. Problems with the Current Definition of Sector Alerts

Volpe has discussed with a number of traffic managers in different FAA facilities the way that TFMS determines if a sector is alerted. The traffic managers have criticized the current method for being a flawed measure of workload, for being overly sensitive, and for focusing on arbitrary, 15-minute intervals.

Flawed measure of workload. The first problem with the current concept is that it does not track well with intuitive notions of sector workload. If, for example, in an interval demand exceeds the MAP (or capacity) threshold for only one minute, with the demand for the other 14 minutes being below the MAP threshold, TFMS will alert the sector for the entire 15-minute interval regardless of the magnitude of the one-minute peak demand. Since ATC might well be able to deal with a problem that lasts for only a minute or two, it seems overly sensitive to issue an alert in this situation without taking into account the magnitude or duration of the excess demand. Moreover, TFMS does not distinguish a congestion problem that lasts for a much longer time (long-term congestion).

Overly sensitive. A second problem with the current concept is that it is an unstable measure that leads to flickering of alerts. For example, often it is the case that when the predicted one-minute peak demand is close to capacity, a small change in the demand (for example, one or two flights) from one traffic update to the next, which is not operationally significant, will cause an alert to go on or to go off. Since Monitor/Alert is updated every minute, tiny changes in predicted demand can cause a sector's alert status to vibrate back and forth, thus pointlessly increasing workload for the traffic manager.

Arbitrariness due to 15-minute boundaries. A third problem with the current concept is that it depends on the accident of where the 15-minute boundaries fall on the timeline. While such intervals might make it easier to present bar graphs, they are not directly related to the traffic in the sky. Desired is a concept that uses the pattern of demand to determine if there is an alert and that is unaffected by artificial boundaries in the timeline.

3. Logic Underlying the Proposed Approach

In preparing this paper we have talked with experienced traffic managers about what intuitively should constitute an alert. In short, an alert should indicate a situation that might lead to a problem and that, therefore, a traffic manager should investigate and decide if any action is needed. The question is: *How can the automation identify a situation that is serious enough to be considered to be an alert?*

What has emerged from these discussions is that *duration* and *magnitude* of excess demand both feed into whether an alert should be declared. Therefore, this report proposes to replaces the current TFMS alert with two kinds of alerts that are called *long-term alerts* and *short-term alerts*. (We are not completely happy with these terms and welcome suggestions.) Section 4 will propose specific rules for deciding when a long-term or short-term alert exists. Before going into these proposed rules in detail, however, we will provide the motivation for this proposal.

Consider what is here called a long-term alert. What led to this idea was traffic managers saying that if demand for a sector is one flight above the MAP value for one minute only and otherwise does not exceed the MAP value, this isolated minute of excess demand should not be enough to trigger an alert. This led to the idea of looking not at a single minute but rather at the pattern of demand over time; an alert is declared only if the excess demand lasts long enough to be of operational significance. Looking not at a single minute but at the pattern of demand over time holds out the possibility of dealing with all three problems with the current concept that were discussed in Section 2. First, the pattern that constitutes an alert can be chosen to try to capture scenarios that represent real traffic management problems. Second, since the pattern takes into account multiple minutes rather than just one, it improves stability. Third, the pattern can ignore the boundaries of 15-minute intervals.

Consider the following example. Suppose that it is decided that the pattern chosen to represent an alert is three consecutive minutes with predicted demand exceeding the MAP. Consider how this would potentially deal with the three problems described in Section 2. First, since three consecutive minutes of excess demand are more likely to represent a true traffic management problem than one minute of excess demand, this pattern is less likely to flag harmless situations as an alert. Second, since this pattern does not allow a single minute of excess demand to trigger an alert, stability is improved. Third, since all that matters is that the MAP is exceeded for three consecutive minutes, it doesn't matter if these three minutes happen to fall into two different 15-minute intervals; therefore, the arbitrary division into 15-minute intervals no longer matters.

Now consider what is here called a short-term alert. While traffic managers report that they typically are not concerned about a single minute that is only one flight above the MAP, they are concerned if the demand for a single minute is considerably above the MAP. Therefore, it should be possible for a single minute to cause an alert if the magnitude of the excess demand for that minute is large enough.

The thinking is that a long-term alert would capture alerts in a way that avoids the three problems discussed in Section 2. However, the long-term alert would miss a situation in which there was very large excess demand for one minute. Therefore, the idea is that by using both long-term and short-term alerts together, traffic managers will get a complete picture of the situations that deserve to be declared as alerts. What is needed is a flexible approach of defining these alerts so that they can be tailored to meet the needs of traffic managers. The next section proposes such an approach for alerting sectors.

4. Proposed New Approach for Declaring Sector Alerts

Two basic questions need to be answered when defining a scheme for determining alerts.

- What causes an alert to be declared?
- What time interval is covered by the alert?

With the current scheme, any minute with demand exceeding the threshold causes an alert to be declared, and the time interval covered by that alert is the entire 15-minute interval containing that minute.

The new proposal uses the demand not for a single minute but rather the pattern of demand over a sequence of minutes to determine if there is an alert and, if so, what time interval is considered to be alerted.

4.1 Long-term Alert

Some terminology is needed here. Any minute for which the predicted demand exceeds the MAP is called an excess or overloaded minute; otherwise, a minute is called a non-overloaded or normal minute. The next few paragraphs discuss patterns of overloaded minutes and normal minutes. An alert will be declared for some pattern of overloaded and normal minutes.

The first question is: What is the minimum number of overloaded minutes that is needed to constitute an *alert*? Let a parameter **a** stand for the minimum number of overloaded minutes sufficient for declaring a long-term alert. In the current TFMS scheme, $\mathbf{a} = 1$. While this is a parameter to be chosen later, it seems likely that it would be greater than 1, e.g., 2, or 3 or 4 or perhaps more.

For concreteness, for the moment set $\mathbf{a} = 3$, i.e., at least three overloaded minutes are required to trigger a sector alert. It is not, however, required that all three overloaded minutes be consecutive. For example, if there are two overloaded minutes followed by one normal minute, and if this pattern repeats indefinitely, there would never be an alert if three consecutive overloaded minutes were required, even though intuitively this seems to be a situation that deserves to be alerted.

Therefore, a second parameter **b** is used to indicate a minimum number of consecutive normal minutes required to "reset" an alert ($\mathbf{b} \ge 1$). This parameter helps determine the start and the end of an alerted interval. This parameter answers the question: *How dense do the overloaded minutes need to be to constitute an alert?*

The proposed rule for a long-term alert can be stated as follows.

- 1. Starting at any minute, count through the minutes, one at a time, and accumulate a count of overloaded minutes along with the counts of consecutive normal minutes separating two adjacent overloaded minutes.
- 2. Whenever consecutive **b** normal minutes are encountered, restart the count at the end of the string of normal minutes.

3. If **a** or more overloaded minutes are encountered before the count restarts, an alert is declared that begins with the first of these overloaded minutes and extends through the last overloaded minute that is followed by **b** normal minutes.

To illustrate how these alerting rules would lead to a long-term alert, consider the example pictured in Figure 4-1. This figure shows predicted sector demand at each minute of the 31-minute period that starts at 1200 and runs through 1230. Assume $\mathbf{a} = 3$, $\mathbf{b} = 4$. This means, roughly speaking, that a long-term alert is declared for any interval with at least three overloaded minutes and with less than four consecutive normal minutes.



Figure 4-1: Example of a Pattern of Demand that Results in Two Intervals of Long-term Alert

In Figure 4-1, the first overloaded minute starts at 1204. Just after 1204, there are three consecutive overloaded minutes followed by five consecutive normal minutes. Hence, there is a demand pattern, which consists of three consecutive overloaded minutes from 1204 through 1206 (with no normal minutes between them) and five normal minutes from 1207 to 1212. Since the alerting rule requires at least $\mathbf{a} = 3$ overloaded minutes for a long-term alert and at least $\mathbf{b} = 4$ consecutive normal minutes to "reset" the alert count, the interval 1204 through 1206 should be alerted. This alert ends after 1206 since there are five normal minutes immediately following 1206.

After the series of five normal minutes, the next overloaded minute is at 1212, so this is the candidate for the beginning of the next alerted interval. Following 1212, the next four consecutive normal minutes that reset the count start at 1223. Between 1212 and 1223 there are five overloaded minutes; since this exceeds $\mathbf{a} = 3$, this means that there is an alert from 1212 through 1222.

It is worth noting that in the pictured interval there are two alerted intervals lasting 3 and 11 minutes for a total of 14 alerted minutes; in contrast, TFMS currently would alert two consecutive 15-minute intervals, i.e., it would alert for 30 minutes from 1200 to 1230.

4.2 Short-term Alert

In our discussions with traffic managers, we have found consensus that the concept of a long-term alert just described is an improvement over the current method of determining an alert. Any alert that the proposed rules pick up does fit their intuition of a real alert. They say, however, that there is an additional class of alerts that these rules do not pick up, and the desire to identify this second class of alerts is the reason for defining what is called a short-term alert.

The thinking of the traffic managers is that if demand for a single minute exceeds the MAP by one flight, this alone is not enough to cause an alert. If, however, the demand greatly exceeds the MAP even for just one minute, this is considered to be significant overloading, and this alone is enough to cause an alert. Since this alert can be triggered by a significant excess demand over the MAP for a short time interval, it is called a short-term alert.

Some terminology is introduced in this section. Let Δ be a parameter that indicates how far above the MAP a higher threshold should be established so that a single minute demand exceeding this threshold indicates that a sector is significantly overloaded. This new, higher threshold is called a short-term MAP and is denoted as MAP_{\Delta} = MAP + \Delta. The assumption is that Δ is set so that a one-minute demand that exceeds the MAP_{\Delta} creates significant sector overload. For example, if we set $\Delta = 3$ and the MAP for a sector is 10, MAP_{\Delta} = 10 + 3 = 13, and a demand of 14 in any one minute alone is enough to cause an alert to be declared. This is called a short-term alert.

It is easy to tell when there is a short-term alert since it is easy to spot a minute when demand exceeds MAP_{Δ} . What calls for some care, however, is to determine the exact interval that is covered by the short-term alert. For example, if a minute in which demand exceeds MAP_{Δ} is followed by a minute that exceeds MAP but not MAP_{Δ} , and if demand for all of the surrounding minutes is below MAP, then what interval does the short-term alert cover?

The proposal is again to use **b** consecutive normal minutes as the "bookends" on each side of an alerted interval. The proposed rule for a short-term alert is:.

- 1. Start with a minute in which demand exceeds MAP_{Δ} . (This minute will definitely be part of an alert. What remains to be seen is how far the alert extends before it and after it and also whether it is part of a short-term or a long-term alert.)
- 2. Go backward in time from that minute until **b** consecutive normal minutes are found. The alert begins immediately after these four normal minutes.
- 3. Go forward in time from that minute until **b** consecutive normal minutes are found. The alert ends immediately after these four normal minutes. This completes the definition of the interval that the alert covers.
- 4. If this alert does not satisfy the definition of a long-term alert, it is a short-term alert.

Short-term alerts are illustrated in Figure 4-2. For this example, set $\mathbf{a} = 3$, $\mathbf{b} = 4$, MAP = 12, and $\Delta = 3$. The short-term MAP, hence, is equal to MAP_{Δ} = 12 + 3 = 15.

Appendix A covers the question of how long a short-term alert can last.



Figure 4-2: Example of a Pattern of Demand that Results in Two Intervals of Short-term Alert

If the pattern of demand is as shown in Figure 4-2, there are two short term alerts. The first is from 1204 through 1208. This is a short term alert because the total number of alerted minutes (two minutes) in this interval is one minute short of the $\mathbf{a} = 3$ overloaded minutes necessary for a long-term alert, and in 1204 the demand exceeds the short-term MAP_{Δ} = 15. The alert begins with 1204 since demand in the four preceding minutes does not exceed MAP and ends after 1208, because there are four consecutive minutes after 1208 where demand does not exceed MAP.

Immediately after the fourth normal minute, there are two consecutive overloaded minutes (shown in Figure 4-2 with bolded bars), followed by four normal minutes. Those two overloaded minutes do not generate a short-term alert because in neither case does demand exceed MAP_{Δ} , and they do not generate a long-term alert since two isolated minutes are not enough to generate a long-term alert. Therefore, even though demand in both minutes exceeds MAP, there is no alert at all for the minutes 1213-1214.

There is a second short-term alert that covers only the single minute of 1219. This is a short-term alert because demand in 1219 exceeds MAP_{Δ} and because immediately on each side of this minute, there are four consecutive normal minutes with demand less than MAP.

Another candidate interval for an alert is the 1224 -1226 interval with two overloaded minutes and one normal minute between them; it is followed by four consecutive normal minutes. As none of the overloaded minutes exceed MAP_{Δ}, there is no short-term alert and the interval 1224 – 1226 is not alerted.

In summary, both of the alerted intervals in Figure 4-2 are short-term alerts because each of them contains less than $\mathbf{a} = 3$ overloaded minutes with less than $\mathbf{b} = 4$ normal minutes between them, and there are enough normal minutes between the alerted intervals to "reset" alerts.

Note: In this paper, for the sake of simplicity, the traffic demands in overloaded minutes were shown by the bars with a single color red. In reality, TFMS shows demand bars with two subsets of traffic demand depending on the flights status. Active flights are shown in the red segments of the bars, and the flights that are still on the ground at the origin airports are shown in the yellow segments of the bars. This paper focuses only on the concepts and will not go into user interface issues such as whether long-term and short-term alerts should be displayed differently.

5. Summary and Next Steps

A view frequently taken by traffic managers is that one minute of excess demand should not constitute an alert unless the excess demand is unusually large. To take account of this view, this report shows how taking account of both the amount of excess demand and its duration can lead to what appears to be a superior method of determining when there is an alert. The proposed method includes parameters that can be adjusted to trade off sensitivity and stability. This approach maintains that, while an alert should not be declared if the demand for one minute is slightly above MAP, an alert should be declared if there is bunching of overloaded minutes. This is called a long-term alert. The parameters **a** and **b** are used to define bunching, where, roughly speaking, **a** indicates how many overloaded minutes are needed to constitute a problem and **b** indicates how concentrated in time these overloaded minutes must be. For example, if one sets $\mathbf{a} = 1$, the proposed method is almost the same as the method currently used by TFMS since it declares an alert if only one minute is even one flight over the MAP. By setting **a** to higher values, sensitivity is decreased and stability is increased. Moreover, the parameter **b** can also be adjusted to further control sensitivity and stability.

The magnitude of excess demand is important since if even a single minute is overloaded enough, this single minute is a problem and justifies an alert for such a short interval. This gives rise to the concept of the short-term alert. The parameter Δ indicates how far above MAP the higher threshold should be so that the demand exceeding this higher threshold will constitute a short-term alert.

The proposed approach holds out the promise of being superior to the current approach used in TFMS since it takes into account both duration and magnitude of excess demand. In particular, the proposed approach is more in line with intuitive notions of excess workload, less prone to instability and flickering, and independent of the arbitrary boundaries between 15-minute intervals. Moreover, there is the promise of more flexibility since there are three parameters that can be adjusted to detect real traffic management problems and to achieve the desired stability of alerts.

As for the next steps in the development of this proposal, if the FAA desires, Volpe could prototype these rules on a test version of ETMS. This test system could be put in front of traffic managers who could both evaluate the general concept of this proposal and also investigate how to set the parameters to reflect their intuition of what constitutes a meaningful alert.

The following questions are among those that could be answered by a prototype.

- 1. Is the concept of a long-term alert as defined above a valid concept that should be operationally deployed? If so, this leads to additional questions.
 - a. What values should be chosen for the parameters **a** and **b**?
 - b. Should the parameters have the same value for every sector in the entire country? Or should each Center be allowed to set the parameters so that they are the same for every sector in the Center? Or should the system allow the parameters to be set differently for every sector, just as MAP values are currently set separately for every sector?
- 2. If the general concept of a long-term alert is valid but there are questions about the particular approach described in this paper, is there some approach that is more useful? For example, would using an average of demand have any advantages?
- 3. Is the concept of the short-term alert as defined above a valid concept that should be operationally deployed? If so, this lead to additional questions.
 - a. What value should be chosen for the parameter Δ ?
 - b. Should this parameter have the same value for every sector in the entire country? Or should each Center be allowed to set the parameter so that it is the same for every sector in the Center? Or

should the system allow the parameter to be set differently for every sector, just as MAP values are currently set separately for every sector?

- 4. If the general concept of a short-term alert is valid but there are questions about the particular approach described in this paper, is there some approach that is more useful?
- 5. Is there any aspect of the current method of declaring alerts that should be saved? For example, if there are minutes for which demand exceeds MAP but which are not enough to cause an alert to be declared, should this be brought to the attention of the traffic manager in some way? See Appendix B for discussion of a possibility.
- 6. If long- and/or short-term alerts are deployed, what should the user interface look like?
 - a. The map display of the TSD indicates alerted sectors with red and yellow polygons with different fill patterns. Should this be changed so that long-term and short-term alerts are displayed differently?
 - b. How should long-term and short-term alerts be displayed in the NAS Monitor? Should this display distinguish long-term alerts from short-term alerts?
 - c. Should there be a new minute-by-minute display similar to the figures shown in this paper?
 - d. Should the user be allowed to toggle between the new approach and the current approach?
 - e. What other aspects of the user interface would need to be adjusted?

Appendix A Duration Limits for Short- and Long-term Alerts

As two new concepts of long- and short-term alerts are introduced, this leads to the question: What is the minimum and maximum possible duration of each of those alerts?

There is a straightforward answer concerning long-term alerts. The minimum possible duration of a long-term alert is \mathbf{a} minutes, when there is no normal minutes between overloaded minutes, and \mathbf{a} is the minimum number of overloaded minutes that constitute long-term alerts. There is no limit for maximum duration of a long-term alert.

As for short-term alerts, only the minimum possible duration can be immediately found: it is one minute. Below, we show how to determine the maximum possible duration of short-term alert τ_{max} .

According to the definition of short-term alert, the maximum possible number of overloaded minutes in a short-term alert is $(\mathbf{a} - 1)$, and maximum number of normal minutes between two closest overloaded minutes that will not cause the "reset" of alert is $(\mathbf{b} - 1)$. To calculate the total maximum possible number of normal minutes in a short-term alert, one should first determine the total number of possible breaks between consecutive overloaded minutes that could be filled with normal minutes. If there are $(\mathbf{a} - 1)$ overloaded minutes, the total number of breaks between them is equal to the number of overloaded minutes minutes one (because one interval is between two overloaded minutes). Hence, there are $(\mathbf{a} - 2)$ breaks between $(\mathbf{a} - 1)$ overloaded minutes. As each break may include maximum $(\mathbf{b} - 1)$ normal minutes, the total maximum possible number of normal minutes in the short-term alert is equal to $(\mathbf{a} - 2)^*(\mathbf{b} - 1)$. Therefore, the maximum duration of a short-term alert $\mathbf{\tau}_{max}$ is equal to the sum of the numbers of overloaded and normal minutes:

$$\tau_{\text{max}} = (\mathbf{a} - 1) + (\mathbf{a} - 2)^*(\mathbf{b} - 1).$$

Figure A-1 illustrates the calculation of maximum duration of a short-term alert in the case of $\mathbf{a} = 4$ and $\mathbf{b} = 4$, with maximum three overloaded minutes $(\mathbf{a} - 1 = 3)$ and with maximum three normal minutes between two consecutive overloaded minutes $(\mathbf{b} - 1 = 3)$. In this Figure, the traffic demands in overloaded and normal minutes are shown by red and green bars, respectively. There are total two breaks dividing the overloaded minutes $(\mathbf{a} - 2 = 2)$, each of which is filled with $(\mathbf{b} - 1) = 3$ normal minutes. As a result, the maximum duration of the short-term alert is equal to $\tau_{max} = (\mathbf{a} - 1) + (\mathbf{a} - 2)^*(\mathbf{b} - 1) = 3 + 2 * 3 = 9$ minutes, and the alert starts at 1205 and ends at the beginning of 1215 (see Figure A- 1).



Figure A-1: Illustration to maximum duration of a short-term alert

In another example with $\mathbf{a} = \mathbf{3}$ and $\mathbf{b} = \mathbf{4}$, the maximum duration of a short-term alert τ_{max} is equal to 5 minutes ($\tau_{max} = (\mathbf{a} - 1) + (\mathbf{a} - 2)^*(\mathbf{b} - 1) = 2 + 1 * 3 = 5$). In this case, the short-term congestion can comprise one or two overloaded minutes, and, can last from one to five minutes. For example, the 5-minute short-term congestion has two overloaded minutes separated by three normal minutes (see Figure A- 2.



Figure A- 2: Five-minute maximum duration of a short-term alert

Appendix B An Option for Displaying Short-term Alerts: Showing Insignificant Overload

According to Section 4.2, a short-term alert is declared for a time interval that contains at least one significantly overloaded minute where demand exceeds MAP_{Δ} . A short-term interval is not alerted if it contains insignificantly overloaded minutes (when one-minute demands exceed MAP but do not exceed MAP_{Δ}).

However, some TFM specialists suggested that they might be interested in viewing such insignificant, short-term sector overload for informational purpose. They would like to have an option of being notified through Monitor/Alert about the insignificant overload. For this purpose, the short-term intervals with insignificantly overloaded minutes can be shown by special coloring of demand bars, e.g., by using a pink color. In this case, the bar chart shown in Figure 4-2 would look like the one in Figure B- 1.



Figure B- 1: Example of Patterns of Demand that Result in the Intervals with Short-term Alert (red) and Insignificant Short-term Overloads (pink)

Figure B-1 displays the insignificantly overloaded intervals (from 1213 to 1215 and from 1224 to 1227) in pink.

For practical purposes, the implementation of new sector alert rules might include an enable/disable option for showing/not showing insignificant short-term sector overloads. That option would allow the TFM specialists to view these alerts when they felt it was necessary and turn them off if they felt it was not needed. It is worth noticing that, from a user standpoint, over time the usefulness of this feature can be reevaluated through feedback from TFM specialists.

Appendix C The Unified Sector Alert Algorithm

Section 4.1 provides the algorithm for detecting a long term alert, and Section 4.2 provides an algorithm for detecting a short-term alert. This appendix gives a single, unified algorithm for alerting sectors that covers both short- and long-term alerts.

- 1. Starting at any minute, look through the minutes, one at a time.
- 2. Find the first overloaded minute and accumulate a count of overloaded minutes along with the counts of consecutive normal minutes separating two adjacent overloaded minutes. Check whether the demand at each overloaded minute is greater than MAP or MAP_{Δ} .
- 3. When consecutive **b** normal minutes are encountered, check the total number of overloaded minutes, including the last one that is followed by the **b** normal minutes.
- 4. If the total number of overloaded minutes is equal or greater than **a** minutes, then a long-term alert is declared that begins with the first of these overloaded minutes and extends to the last overloaded minute that is followed by **b** normal minutes. Then restart the counts from step 2.
- 5. If the total number of overloaded minutes is less than **a** minutes but if the demand for any minute in this interval exceeds MAP_{Δ} , then a short-term alert is declared for the whole interval that includes the first and the last overloaded minutes. Then restart the counts from step 2.