

TSP & Preemption

Introduction, Basic Concepts and Guidance



SFMTA

This memo is intended to provide background and guidelines for the implementation of Transit Signal Priority (TSP) and preemption for transit, emergency vehicles, and other situations (e.g., draw bridges, heavy rail) in San Francisco. This memo addresses TSP for buses, light rail vehicles (LRVs) and streetcars operating in mixed-flow traffic lanes and on corridors with transit-only lanes or dedicated right-of-way.

When properly implemented, TSP can be one of the most cost-effective tools to reduce travel times, reduce delay, and improve reliability along a transit line. Like any tool, it must be used correctly and in the right circumstances in order to be effective. An understanding of how TSP works and its limitations is critical in ensuring success. The guidance provided in this document is meant to help staff make informed decisions when deploying and modifying TSP. This document is a guide to the more common TSP features and settings, and is not a substitute for engineering judgment. For more detailed information and guidance on specific applications, refer to the D4 manual and/or consult with the ITS-SFgo team for current practices. Exceptions to this guidance may be appropriate on a case-by-case basis.

Recommended Parameters: Page by Page

The guidance regarding parameters for each Timing Card page follows each section in a box similar to this one. Engineers should also refer to the “TSP Settings Explained” example timing card spreadsheet.

Table of Contents

1.	TSP & Preemption Fundamentals	4
1.1.	What is TSP?.....	4
1.2.	What is the difference between TSP and Preemption?.....	4
1.3.	What about passive TSP?	4
2.	TSP Basics	5
2.1.	Determining time available for TSP: Cycle Float	5
2.2.	Green Extension.....	6
2.3.	Early Green	7
2.4.	Changing or inserting phases with TSP or Preemption	8
2.5.	Can you use TSP with nearside transit stops?	9
2.6.	Can you provide Conditional TSP or Preempt based on certain conditions?	9
2.7.	What happens if more than one TSP call comes in at the same time?	9
3.	TSP – General page	10
3.1.	Priority Minimum Green (used for EARLY GREEN)	10
3.2.	Recovery Minimum Green (Used for EXTENSIONS).....	10
3.3.	Drop Free	10
3.4.	Priority Alternate Sequence (Phase Swapping)	11
4.	Transit Timing	12
4.1.	Coordination Priority Modes	12
4.2.	Coordination Extension Limit:	12
4.3.	Reservice Inhibit.....	12
4.4.	TSP and Free Operation.....	12
4.5.	Call mode	13
4.6.	Rail / Dedicated Transit Phase Settings.....	14
5.	Transit (or Emergency Vehicle) Detection	16
5.1.	General Detector Settings:	16
5.2.	Vetag-Specific Settings:	17
5.3.	Other Settings on Transit Detector Page	18
6.	Preemption (Transit or Emergency Vehicle).....	19
6.1.	Preempt Settings	20
6.2.	Getting Back in Step - Custom Transitioning.....	20

7.	APPENDIX - Advanced Features Rarely or No Longer Used	24
7.1.	Travel Time Slack	24
7.2.	Estimated Delay	24
7.3.	Adaptive Priority	25

1. TSP & Preemption Fundamentals

1.1. What is TSP?

TSP prioritizes transit vehicles through a signalized intersection. Transit vehicles communicate their position via radios/transponders to detection hardware connected to an intersection's traffic signal controller. (Technology is evolving, and there are cloud-based solutions in development that may allow for real-time communication of the transit vehicle's location to the traffic controller without dedicated hardware.) The signal controller then considers the TSP parameters in its program, and determines if and how to modify the signal timing to minimize signal delay for the approaching transit vehicle.

Specifically, TSP can:

- Extend the transit vehicle's green light
- Bring the transit vehicle's green up earlier by shortening the opposing red light.
- Rearrange signal phases to reduce the signal delay for transit

1.2. What is the difference between TSP and Preemption?

Preempts more aggressively prioritize transit, but are more disruptive to normal traffic operations than TSP. Preempts are reserved for the special circumstances, such as cable car operations (limited ability to brake/stop), railroad crossings and drawbridges, emergency vehicles with active sirens, unusual transit maneuvers or actuated transit phases that cannot be accommodated in the traffic cycle.

Preempts are different from TSP because they can:

- truncate (i.e. immediately end) Walk and/or Flashing Red Hand pedestrian phases.
- Insert special phases that can't be fit within the typical signal cycle
- Hold the desired phase for extended periods.

Compared against preempts, TSP can only modestly extend, truncate, or re-order phases because the settings limit it getting back into coordination ("in step") within one signal cycle. TSP typically uses priority and recovery minimum timings that are longer than the absolute minimum timings used by preempts. Because preempts are allowed to make more significant changes to the signal timing, a preempt could require several subsequent signal cycles to get back in step.

1.3. What about passive TSP?

This memo focuses on traditional "active" TSP, not "passive" TSP. Passive TSP uses fixed signal timing strategies to prioritize transit vehicles without any real-time communication regarding the transit vehicle's location. A passive priority timing plan considers stop locations, average dwell times and transit travel times to inform the corridor's cycle length, phase splits, and offsets to provide progression for transit. Passive priority should be considered on corridors with frequent and predictable transit operations, though many factors may limit the feasibility of implementation. Active TSP can work in conjunction with passive priority to further improve transit travel time and reliability through signalized intersections. Van Ness BRT is an example of a corridor with both passive and active TSP. For the remainder of this memo, references to "TSP" refer to active TSP.

2. TSP Basics

2.1. Determining time available for TSP: Cycle Float


Before you can determine how much TSP can be provided, you need to know the available “float”. The float within the traffic signal cycle is defined as the difference between the selected cycle length (e.g. 60, 75, or 90 seconds) and the minimum cycle length. Minimum cycle length is the shortest duration that the signal cycle could time assuming that every vehicle and pedestrian phase is actuated and displays for its minimum duration. The float for each phase is the “extra” time that the signal phase receives in excess of the minimum timing, which is typically driven by the pedestrian signal timing calculations. For left turn phases, engineering judgement should be used to determine minimum green values for TSP applications.

For intersections with pedestrian phases, float is limited to any extra time in the “Green/Walk” time above the minimum WALK. If there are no pedestrian signals, the float should still be calculated such that the minimum split would provide enough time to cross the roadway. If there are signalized left turn phases, engineering judgement should be used to determine minimum green values and if any float can be taken for TSP.

The cycle float is the sum of each phase’s float in the cycle. Because the cycle float is the total amount of time that can be skipped in cycle without violating minimums, it represents the amount of time that can be shifted around to accommodate an approaching transit vehicle for TSP. Limiting the maximum early green and the maximum green extension to the cycle float allows the controller to step back into coordination within one cycle. The float could come from the main street, the side street, or a combination of both. In some cases it may not be desirable to use the maximum available float; the engineer may want to restrict the length of the extensions or to ensure that a sensitive phase isn’t shortened to its absolute minimum for capacity or coordination. The float is not input anywhere in the controller; it is just information for the engineer to use when setting the **minimum priority greens** and **minimum recovery greens**, which informs the controller which phases can be shortened to provide TSP.

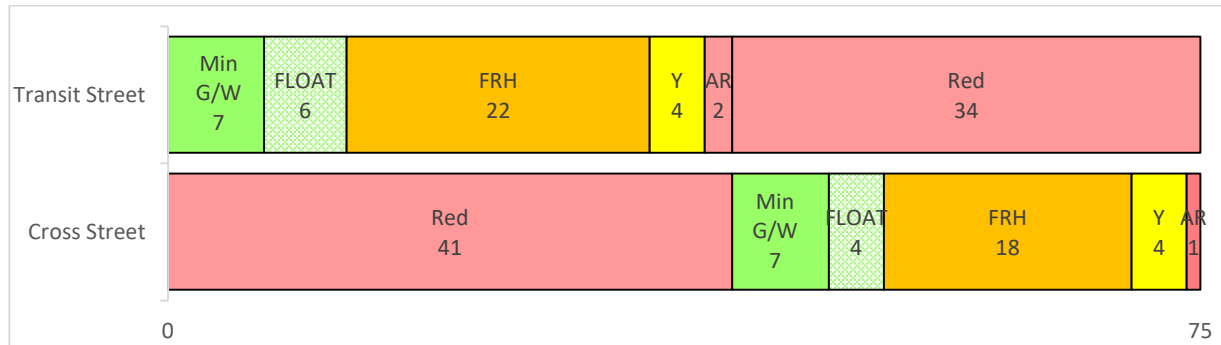
Example: The transit street has a split of 36 seconds. Based on the minimum timings, there is 10 seconds of float available.

Figure 1. Float is the extra time available after accounting for minimum timings.



Example: The transit street has 6 seconds of float, and the cross street has 4 seconds of float. The total cycle float is 10 seconds. The engineer could determine if the full 10 seconds of float should be used, or set the minimum and recovery priority greens to indicate which phases can be shortened.

Figure 2. Determining cycle float from multiple phases: 6 sec from transit street and 4 sec from cross streets = 10 seconds total float



The traffic controller uses the available cycle float to adjust the signal timing to reduce transit delays. When the intersection traffic signal controller receives a TSP request, the controller will use the expected time to arrival and the current point in the cycle to decide which action to take. Note that traffic controller may have green extension and/or early green enabled; the type of TSP that is allowed is determined by which Priority Coordination mode is selected (see section 4.1).

2.2. Green Extension

The controller will provide a green extension if it determines that the vehicle will reach the intersection within the duration of a possible green extension. After a green extension, the controller enters a recovery cycle to recuperate the time borrowed for an extension. The recovery cycle allows the controller to get back into coordination by taking float time wherever available, which may include shortening the green for the side streets, turn phases, or the transit street as allowed.

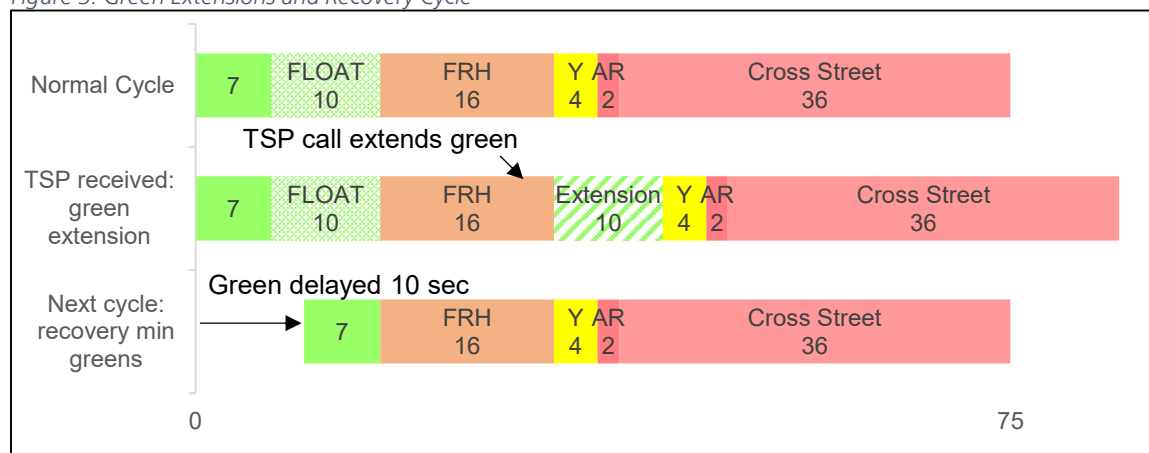
- *Minimum extension:* Operators may see pedestrian countdown signals approaching zero as they approach an intersection and slow down anticipating the start of yellow, unaware that the traffic controller is preparing extend the green light with a RED HAND display for pedestrians. Thus, we use a minimum extension of 7 seconds to give transit vehicle operators sufficient time to notice that a green extension has been granted and take advantage of the extension. If there is not enough cycle float to provide this minimum value, TSP should not be implemented. The minimum extension is not explicitly programmed into the D4 program.
- *Maximum Extension:* the length of the green extension is limited to reduce the effects of excessively long red intervals for other phases, which could lead to operational concerns. For this reason, we recommend limiting extensions to 15 seconds for cycles that are 75 seconds or shorter, or 20 seconds for cycles longer than 75 seconds. The green will only be extended until the transit vehicle has passed through the intersection; the maximum TSP extension time is only fully used if it is needed.
- Different extension lengths can be input by dial; TSP can also be disabled by approach and/or by dial.

- If the available TSP extension time is calculated to be insufficient for the transit vehicle to get through on green, the controller will not provide an extension.

Cycle Length	Recommended Max Extension
75 seconds or less	7 to 15 seconds
Over 75 seconds	7 to 20 seconds

Example – granting green extension: Building on the example in Figure 1, there is 10 seconds of float. Thus, the possible extension duration is 10 seconds. If a TSP call comes in just before the start of yellow with a travel time of 8 seconds, the controller will grant an extension. The extension will end once the bus drops its call (after it passes through the intersection) or after 10 seconds, whichever is first.

Figure 3. Green Extensions and Recovery Cycle



Example – NOT granting a green extension: A bus is approaching just as the transit phase is about to turn yellow. The estimated travel time is 12 seconds, but the available extension is only 10 seconds. The controller will not provide an extension.

2.3. Early Green

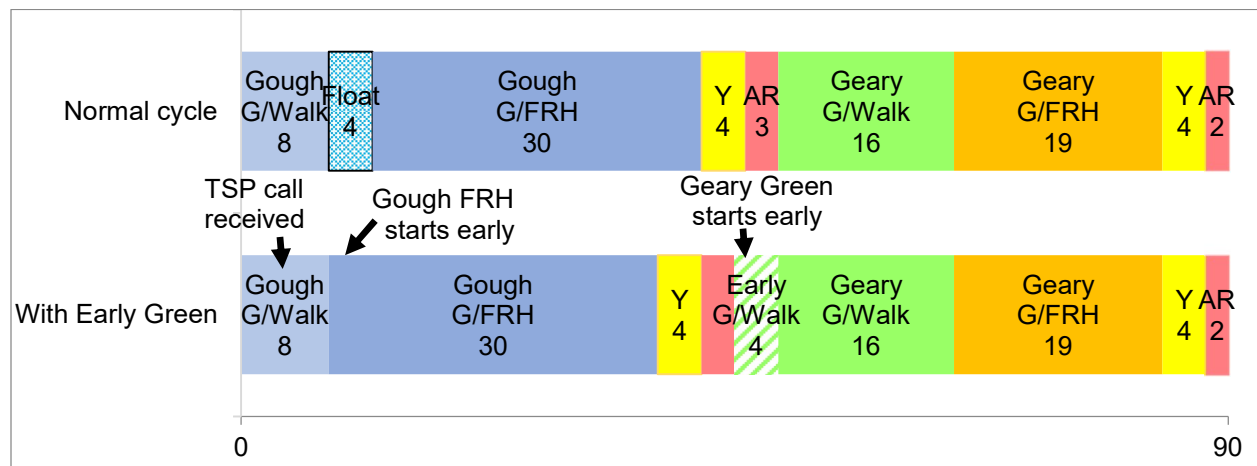
If the transit vehicle is expected to arrive after the transit phase turns red, the controller could bring up the next transit phase green sooner by shortening conflicting phases. The cycle float taken from the shortened phases will be appended to the next transit green.

Early green is most effective if there is float available from the cross street AND pedestrian phases are actuated, or if there is significant advance notice of the transit vehicle’s arrival. Otherwise, early green is quite limited in its capabilities. Since most intersections have pedestrian phases on recall, any call for early green would have to be received early enough in the WALK phase to be able to start FRH early.

Early TSP is less disruptive to coordination than extensions because the beginning of green is not delayed, something that is more noticeable on a coordinated arterial such as Gough or Franklin. For example, the signal at Geary and Gough is programmed such that Geary bus that arrives during Gough’s WALK/Green phase can shorten the Gough green by up to 4 seconds. The Geary green will

display 4 seconds earlier, and the Geary phase will be 4 seconds longer. The following cycle, the Gough phase will start of green at the same time, therefore coordination isn't disrupted - only the end of the green.

Example: Geary Boulevard and Gough Street: Engineer decides Geary should not be shortened but Gough Green can be shortened up to 4 seconds. By enabling 4 seconds of early green, the controller can take the float from Gough Street and give it to Geary Boulevard. However, the early green can only be granted if the TSP call is received before or during the first 8 seconds of Gough Green, so that the FRH can start early. The start point of Gough Green will be unchanged. However, the transit street green will come up early prior to the typical offset.



2.4. Changing or inserting phases with TSP or Preemption

With San Francisco’s D4 traffic signal software program, you can change the phase order or insert phases using either TSP or Preempt. The key difference is that preempts allow the signal to get out of step for more than one cycle.

2.4.1. Priority Alternate Sequence (TSP)

The priority alternate sequence allows the controller to re-order phases to reduce red delay to the transit vehicle. An example would include switching a leading left into a lagging left so that the main transit street serves prior to the main street turn phase. This application is more commonly known as “phase swapping”. Phase swapping has been used at many locations along T Third, such as 4th/Channel, since 2007. This feature is also used at 3rd/Kirkwood to insert an actuated transit phase either before or after the main street phase.

2.4.2. Phase insertion through Preempt

At locations where the transit phase cannot be inserted within the regular signal cycle or where you want to provide the most flexibility for inserting a phase, a preempt can be used. This is currently used at numerous locations, such as at exit points to the Van Ness BRT busway, or LRV pull-in and pull-out movements near Muni Metro East (MME).

2.5. Can you use TSP with nearside transit stops?

For buses – No. TSP should only be implemented at intersections for bus routes that do not have nearside transit stops for two reasons. First, dwell times can be variable, especially if a passenger requires the ADA ramp or is loading a bicycle on the bike rack. Second, the current technology only sends requests for TSP while the doors are closed. If a bus is approaching a nearside stop, it would request TSP as it approaches the stop, then drop the call while it is loading, then request TSP after it is finished loading. The traffic controller cannot distinguish between a bus approaching the nearside stop and a bus that is finished loading and ready to depart.

For Rail – Yes. TSP can be used at locations with nearside stops as long as there is a vetag detector located between the transit stop and the stop line. LRV dwell times tend to be less variable due to level boarding and multiple doors. Upstream vetag detectors will indicate when the LRV is approaching the transit stop, and the programming can be set to account for the LRV's anticipated dwell time. When an LRV finishes loading and pulls forward to be detected by the stop bar vetag, the signal will get a "check-in" call from the vetag that is specifically designed to be placed beyond the transit stop location.

2.6. Can you provide Conditional TSP or Preempt based on certain conditions?

Conditional TSP is possible, for example, providing priority for a Rapid route, but not a Local route, as long as the Rapid and Local routes are programmed as different classes in the GTFS.¹ For example, the Geary Rapid buses are programmed to be Class 8 and the Geary Local is Class 4. For example, we have enabled transit priority for ONLY the 38R Geary Rapid at Geary/Webster, the 38 Local has a nearside stop, but the Rapid does not stop by indicating that Class 4 should not receive TSP.

We also use Classes to distinguish between movements requesting a preempt. Along Van Ness Avenue, all Golden Gate Transit (GGT) buses can be recognized by the controller because they are Class 7. At Van Ness Avenue and Golden Gate Avenue, GGT buses call for a preempt to get an actuated phase to turn left out of the busway through preempt that is only granted to Class 7 buses. However, at other locations, GGT buses are not treated differently from Muni buses because unless otherwise stated, the default is to provide TSP to all classes.

For further guidance, consult with SFgo staff, as there are limitations on how classes are set up and can be used.

2.7. What happens if more than one TSP call comes in at the same time?

The traffic controller will provide TSP on whichever phase is playing, even if a TSP call for a conflicting phase arrives first. (TSP is programmed for "all active channels", not "first-come first-served".) We can currently disable TSP per approach or per dial, for example if we want TSP in a non-peak direction to be off, however we cannot designate if one direction is a higher priority than the other.

¹ GTFS is "General Transit Feed Specification", an online feed that connects each transit vehicle to a route and schedule for each day, for example, that Bus 7256 is running on the 49 Van Ness route.

3. TSP – General page

The TSP General tab sets the **Priority** and **Recovery Min Greens** to ensure that each phase displays the desired amount of time, if it is not desirable to shorten a phase to its absolute minimum green. These settings let the controller know which phases it can shorten and by how much.

3.1. Priority Minimum Green (used for EARLY GREEN)

The priority minimum green parameter establishes the shortest green time phase when the controller is attempting early TSP; which might be longer than the minimum green on the Base Timing page. Priority minimum greens will influence which phases are shortened (or not shortened) during the early TSP cycle. Engineering judgement should be used in determining whether to use a priority minimum green for a given phase, or if the base timings are sufficient during TSP.

An example would be at locations with high pedestrian or vehicle volumes. Higher pedestrian volumes may lead to queuing to enter the crosswalk and traffic volumes may support longer green times for coordinated corridors. At locations like these, a higher priority minimum green time than the absolute minimum green time (as determined by min walk and FRH) should be used to balance all modes. Note that early TSP is by nature less disruptive to coordination, since it shortens the end of the green for the movement(s) opposing the transit approach.

3.2. Recovery Minimum Green (Used for EXTENSIONS)

The recovery minimum green parameter establishes the minimum green time required to be reserved for a phase in the cycle following the TSP extension cycle. When TSP extensions occurs, the controller will determine minimum green values by selecting between the largest of the recovery minimum green and the base timing minimum green values. Note that LPI values do not count toward the minimum green time. In practice, this parameter limits how long a TSP extension can occur, and helps determine which phase(s) the float will be borrowed from. Time used for TSP extensions can be borrowed from any phase, but is typically borrowed from its own transit phase in the next cycle. Therefore, recovery minimum greens shall be used in all phases with float.

3.3. Drop Free

With Drop Free, the controller will drop out of coordination into Free mode to serve a TSP request. When this occurs, the controller uses the free extension and hold limits to determine the limits of TSP. This is used for locations where there is not enough float in the cycle to provide TSP, but the engineer does not want to increase the cycle length to accommodate TSP and is willing to take the signal out of coordination in order to reduce transit delays. Drop Free will serves a TSP extension without being constrained by the cycle length. (Priority minimum greens apply, but not recovery minimum greens.) Once TSP is served, the controller will return to coordination and use the Base Timing page's Max/Min transition parameters to get in-step. Drop Free is currently used at Bayshore/Hester and at some locations along Haight Street. (There is also an option to use "Drop Free In Sync", which is similar

except that after serving priority the controller will return to the next phase to get in-step without custom transitioning, thus may result in phase skipping. This is not currently used.)

3.4. Priority Alternate Sequence (Phase Swapping)

The priority alternate sequence allows the controller to re-order phases, for example switching a leading left into a lagging left so that the main transit street serves prior to the main street turn phase. This application is more commonly known as “phase swapping”. Early TSP must be enabled for alternate phase sequence to work in D4. For intersections that are manual control plans for Parking Control Officers (PCOs), alternate phase sequence shall not be used since the PCO cannot effectively manage traffic if they cannot predict which phase will display next.

Recommended Parameters – TSP General Settings

- Determine the available float to determine how much you have available to extend / truncate for transit priority calls. Note that if you are taking time from multiple phases, you can distribute the truncation / extension amount to each phase you want to modify as desired.
- **Priority Min Green** (for Early TSP):
 - Use formula: [Total green - Truncation Amount]. **Exclude LPIs***
 - Early TSP is recommended to allow the transit phase to come up early by between 4 seconds (principal arterials) to 10 seconds (minor arterials or lesser streets). However, there is no minimum for enabling Early green.
 - If a phase should not be shortened for a priority call, enter “255”. (The controller will default to the steady demand timing.)
- **Recovery Min Green** (for TSP extensions):
 - Use formula: [Total green - Extension Limit]. **Exclude LPIs***
 - If a phase should not be shortened for recovering from a priority call, enter “255”. (The controller will default to the steady demand timing.)
- **Priority Alternate Sequence** is used in rare, special circumstances, such as intersection with an actuated transit phases that could be inserted in different places in the cycle (e.g. Bayshore/Hester, 3rd/Thomas) or where there is a leading left-turn phase that could be shifted to be lagging (e.g. Embarcadero/Folsom, 3rd/Bayview/Revere). This feature is not effective for lagging left turns. Consult with Transit Engineering staff to determine if it may be appropriate. If used, enter language similar to “Swap phases X and X” in the appropriate dials when phase swapping is allowed. Early TSP must be enabled for alternate phase sequence to operate in D4.
- **Free Priority** - Priority Min Green is typically calculated just like coordinated Priority Min Green; Base Timing Mins are generally adequate.

4. Transit Timing

4.1. Coordination Priority Modes

There are six types of coordination priority modes: None, Early / Extend, Extend Only, Early / Extend Reservice, Drop Free, and Drop Free In Sync. The first three options are the most frequently used, and are defined in the Section 1. Drop Free and Drop Free In Sync are discussed in Sections 3.3.

4.2. Coordination Extension Limit:

Transit priority extensions should not be enabled unless there is 7 seconds of permissible float. The recommended extension length is 7 and 15 seconds if the cycle length is 75 seconds or lower, and between 7 and 20 seconds for cycle lengths greater than 75 seconds. Extension Limit settings can be different by dial, and extensions can be disabled by dial or by approach. The extension limits can vary by dial. (e.g., 15/15/0 means 15 second extensions for Dials 1 and 2, but none for Dial 3.)

4.3. Reservice Inhibit

To maintain coordination, TSP is set up such that a traffic controller could be “back in step” serving the coordinated phases at the correct time with respect to the offset no later than the end of the subsequent cycle after a TSP call is served. If a second TSP call is received shortly after the first one has been served, the “reservice inhibit” settings determine if the traffic controller could “reservice” that transit phase by providing TSP, or it should ignore the transit call and focus on getting back into step. Reservice inhibit can be programmed to prevent a subsequent TSP call from being granted within a specified interval to ensure that the intersection returns to coordination before providing TSP.

Different reservice inhibit settings can be set for transit calls from the same direction versus any transit call. Typically, the reservice inhibit value is set to (cycle), to allow the traffic controller one full cycle to recover before responding to a subsequent TSP call.

Example: If an inbound train is closely followed by another inbound train, you may not want to provide TSP to the second train to reduce bunching. However, you may want to give TSP to an outbound train that arrives shortly after an inbound train, since it is coming from a different direction. **Reservice Inhibit Same** prevents transit signal priority during the inhibit period specified for the same priority direction only (e.g., no priority for second inbound train). If you did not want subsequent TSP calls served regardless of direction, you would use the **Reservice Inhibit All** timer (e.g., both inbound and outbound are inhibited).

If reservice inhibit is not enabled, the signal could continually be out-of-step for multiple cycles when a transit vehicle is present every cycle such as on high volume transit corridors on Mission or Geary Streets.

4.4. TSP and Free Operation

The discussion above refers to when traffic signals are operating in Coordinated mode. However, sometimes traffic signals operate in Free mode, where there are minimums and maximums for each phase, and at least one phase is actuated. Free mode is different than coordinated modes in that it

does not follow a coordinated offset, is not constrained to a defined cycle length, and requires detection. Since there is no defined cycle length, TSP is not constrained by cycle float. However, a potential disadvantage of Free is that there is no coordination of signal progression along a corridor. Therefore, Free operation can result in higher travel times than coordinated signals with TSP. When a three-day test of Free operations was conducted on the T Third in 2013, the travel times increased by 5% from 59 minutes round trip to 61.5 minutes. Thus, using TSP and Free operation is most appropriate for locations where there is not signal progression / coordination or during off-peak times.

4.4.1. Free Hold Limit - NO LONGER USED

The Free Hold Limit applies only to traffic signals that do not operate in Coordinated mode. The green light will only hold when the transit vehicle travel time is equal to or less than the free hold limit parameter. This feature helps limit long side-street delays when a transit vehicle is approaching. Ideally, this parameter can be set equal to or less than the side-street split time. If the transit travel time is greater than the side street split time, the side street can be served without unnecessary delay.

4.4.2. Free Extension Limit

This parameter is the same as the coordinated extension limit, but is used only when the controller is running Free. The free extension limit is similar to the coordinated extension limit: this value is the maximum green extension that a transit vehicle can receive.

4.5. Call mode

This setting designates the call mode of the transit detector for actuated situations. Options include recall, soft recall, lock, and non-lock. Please refer to the D4 manual for details on each call mode.

Recommended Parameters – Transit Timing

- **Coordination Priority Mode:** check the box for **Early/Extend**, **Early only*** or **Extend Only**. The controller will provide early and/or extend as it is able based on the priority and/or recovery min green settings on the “TSP General” tab.
 - Early/Ext Rsv - not used; listed just to match D4 menu of options
 - **Drop Free** – advanced feature used in special cases (described in Section 3).
 - Drop Free IS – not used; listed just to match D4 menu of options
- **Coordination Extension Limit:** For cycles up to 75 seconds, use 7 to 15 seconds. For Cycle >75, use 7 to 20 seconds. If using different extension limits by dial, enter a number for each dial (e.g. 15/15/0 means 15 second extensions for Dials 1 and 2, but none for Dial 3.)
- There is no box or field to enter a discrete value for early TSP limits; it is solely controlled by the Priority Min Green on the TSP General setting page. However, engineers are encouraged to note the intended amount of time early TSP can shorten in the notes section, and which phases early green TSP may shorten.
- **Reservice inhibit** – in mode box, enter “cycle”. (earlier, we also used “seconds”, but this is no longer used.) In the boxes below, enter a number (typically “1”) in EITHER the **Same TSP Request** (a call in the same direction) or **ALL TSP Request** (calls from any direction) **box**. This will indicate how many cycles subsequent TSP requests should be ignored.
- **Free Priority Mode** – typically marked as “None” unless schedule includes Free operations. If used, engineering judgment should be used to determine the **Free Extension Limit**.
- **Call Mode** – typically “recall” or “non-locking” (if actuated).

**While “early only” is a choice on the TSP settings page, this is only for engineering clarification; electricians will actually select Early/Extend.*

4.6. Rail / Dedicated Transit Phase Settings

There are some additional settings for intersection with Rail or other dedicated transit signals.

4.6.1. Advance Call

This parameter indicates the preferred amount of time that a transit phase is green prior to its arrival at the intersection, essentially trying to have the transit phase be green far enough in advance so that it does not have to decelerate as it approaches the signal. This means that if the light is red as the transit vehicle approaches, the controller will try to have the light green for transit sooner, using the calculation [estimated travel time – advance green]. For example, if a train will arrive in 30 seconds, but we want the light green 10 seconds before the LRV arrives, the controller will use an arrival time of 20 seconds if the light is currently red. However, if the light is already green, it will use the full estimated travel time to predict when the train will arrive. The default value for this is 10 seconds.

4.6.1. Advance Call

Once the LRV phase travel timer (estimated time until arrival at the intersection) is equal to or less than this value, a call will be placed to serve the LRV phase when operating without transit priority. When the LRV phase has the priority mode enabled, the Advanced Call Time will be subtracted from the Travel Time to compute the target LRV green time. This parameter attempts to reduce the amount of slowing and stopping to LRV operators by providing a green in advance of the transit vehicle's arrival, to avoid the operator slowing down if they see a red light as they approach the intersection. This feature does not apply when the controller is determining travel time for TSP extensions.

4.6.1. Advanced Warning Time

This parameter sets how many seconds in advance the flashing train coming signs (FTC Signs) will be displayed prior to the transit vehicle arrival. The warning sign will be activated once the estimated travel time is less than or equal to this parameter. If not specified, the signal shop will assume 15 seconds.

4.6.2. Advance Green (Actuated phase only)

This rarely-used parameter is only used for actuated transit phases (those that utilize the Locking Call or Non Locking Call mode), AND can time with parent phases. Once the transit travel timer (estimated travel time to the intersection) is equal to or less than this value, the transit phase will turn green only if the parent phases are green and can stay green for long enough to serve the full transit phase before force-off or max-out. One location where this is used is 19th Avenue and Rossmoor.

4.6.3. Warn mode

This feature defines when the Train Coming sign will turn on. The default setting for this is "G+Ycall" so that the Flashing Train Coming signs will only display when the transit phase is green or yellow and there is transit vehicle detected.

Recommended Parameters – Transit Timing (Rail/Dedicated Transit Phase)

Intersection with dedicated transit signals (e.g. white/red lunar bars) have additional settings

- **Min Green:** typically 10 seconds
- **Advance Call:** typically 10 seconds
- **Call Mode:** "recall" or "non-locking" (for actuation)
- **Yellow:** typically 6 seconds for transit phases (e.g. white lunar bar)
- **Advance Green:** used only for actuation; typically 10 seconds.
- **Warn mode:** enter "G+Ycall" if Flashing Train Coming signs are present. This allows the FTCs to display only when the transit phase is yellow or green and a train is present.
- **Train Coming:** enter the phases that have FTC signs.

5. Transit (or Emergency Vehicle) Detection

TSP relies on detectors that can detect transit vehicles to inform the traffic controller of the expected transit arrival time. Local detection is defined as a detector or detection area for the transit vehicle that is upstream of the signal. There should be no traffic signals, stop signs, or transit stops between the vehicle and the signal that the local detector is activating.

Remote detection enables the tracking of transit vehicles as they approach and pass through multiple-TSP enabled signalized intersections along a corridor. Remote detectors function identically to local detectors by providing information about the transit vehicle's location by relaying the TSP call from an upstream intersections' detectors. Remote detectors provide more advance notice, which allows the controller to make decisions on which TSP method to provide.

Remote Detectors are typically not used in mixed flow conditions due to the increased variability in travel time caused by congestion. On corridors with dedicated right-of-way, signals may provide TSP using up to four remote detectors. The accuracy of the expected arrival time decreases the further away the detectors are from the local intersection due to increasing variability in travel time, transit stops dwell times, traffic controls, etc. as the distance increases. While we initially used remote detectors a full cycle length of travel time away on T Third, we later reduced this to half a cycle to minimize programming effort as there was diminishing returns with that much advance notice. Care should be taken to determine how many remote detectors are appropriate for each application.

Remote detectors are programmed over a peer IP network, thus cannot be used if there is no intersection communication (fiber, network radios).

5.1. General Detector Settings:

5.1.1. Travel Time

This is the estimated travel time from when the transit vehicle first enters the detection zone to the center of the intersection (for both local and remote), assuming that the vehicle experiences no red light delay at intermediate intersections (if any). The travel timer will begin to count down as soon as the transit vehicle triggers the detector. At this time, the various detectors used for transit priority typically alert the traffic signal when a transit vehicle has reached a certain threshold distance away from the intersection, then the controller assumes a preset travel time from that point to the intersection, and attempts to have the signal be green for the transit vehicle when it arrives. Currently, we cannot adjust the estimated travel time based on an individual vehicle's second-by-second movement, so unforeseen delays may result in inaccurate travel timers. This value should typically represent the *free-flow speed at which the transit vehicle travels* and is calculated by dividing the distance by an assumed speed based on engineering judgement. As new data tools become available, we may use AVL data as appropriate.

5.1.2. Travel Time Slack – NOT LONGER USED

The Travel Time Slack feature is no longer recommended. See appendix for additional info.

5.1.3. Allowed late arrival to next detector (Remote Call Max Time)

This value represents how much *additional* time a remote call will be held if it does not reach the next detector by the programmed travel time.

Example: Remote call 1 has a travel time of 30 seconds and the local detector has a travel time of 10 seconds. When the remote call comes in, the time starts at 30 seconds then counts down. When it reaches 10 seconds, the controller would anticipate the vehicle to reach the local detector. Unless there is an “allowed late arrival” value entered, the remote call will be dropped at the 10 second mark if the local detector has not been actuated. However, if there is an “allowed late arrival time” of 5 seconds, the controller will hold the remote detector call for an additional 5 seconds. After the additional 5 seconds, the TSP call will be dropped if the local detector still has not been activated. If a max time of 15 seconds is programmed, the controller will count down the travel timer to 0, then hold the 0 for another 5 seconds. After the full 15s of max time is complete, and there is still no local detector call, the remote call is dropped.

5.1.4. Time added if late (Remote Call Adjust) – NO LONGER USED

This value is the amount of time added to the travel time if the vehicle is late to the next detector. Generally, this setting is no longer used, as it has caused detector calls to be locked, leading to unnecessary TSP extensions.

5.2. Vetag-Specific Settings:

5.2.1. Checkout Limit

The checkout limit timer designates the amount of time after a check-in, but before a check-out that would consider a failed checkout detector. When the timer limit is reached, the controller will automatically checkout the phase and the controller will count a failed detector occurrence.

5.2.2. Checkout Fail Count

The checkout fail count is the threshold that needs to be reached for the controller to place the phase on recall. A checkout failure is counted when there is a check-in without a check-out for the duration of the checkout limit timer.

5.2.3. Checkout Mode

Typical check-out modes are Normal and Count. Count differs from normal in that it counts the number of check-ins and check-outs and does not check-out the phase until the difference of the two are 0.

5.3. Other Settings on Transit Detector Page

These settings are not used for typical TSP installations, but may be appropriate for certain locations or corridors, such as the T-Third route or BRT corridors, where the vehicle is in dedicated Right of Way.

5.3.1. Warning Extension

This setting is used to extend the FTC display time after a check-out is received. The default value for this is 6 seconds.

5.3.2. Adaptive Priority – NO LONGER USED

Adaptive priority dynamically adjusts the travel times by “learning” from previous transit vehicle arrivals. In practice, adaptive priority did not improve TSP effectiveness and in some cases had negative impacts, so D4’s Adaptive Priority settings are no longer used. See the Appendix for additional information.

Recommended Parameters – Transit Detectors

- **Delay, Extend:** Not used, leave blank.
- **Warning Extension:** enter 6 if there are FTCs, otherwise leave blank.
- **Checkout Limit:** Enter 255 for rail only, otherwise leave blank.
- **Checkout Mode:** enter “count” for rail only, otherwise leave blank.
- **Checkout Fail Count:** Enter 1 for rail only, otherwise leave blank.
- **Transit Stop location** – mark the appropriate location with an X (Near-side, midblock, Farside, N/A)

Local Detection zones - Bus TSP should use local detection zones set such that TSP is not activated until the transit vehicle has finished serving any upstream transit stops and passed through any upstream controlled intersections. The distances are measured from the center of the local intersection. In the local detection section, complete the following entries:

- **Intersection Number** - This is referring to the intersection’s peer address on the peer network, which is assigned by the Signal Shop when the peer network is configured. Function -For Bus detection, typically “TSP Call”. For Rail using vetag, “Check-in” or “Check-out” as applicable.
- **Location** – leave blank for local detectors; specify intersection for remote detectors.
- **Type** – For bus, typically “GPS Low Priority”. For rail, typically “vetag” or “VPI” (if vetag calls are controlled by a “Vital Processor Interlocking”, the track switch controller).
- **Travel Time** – Travel Time in seconds from Detector to center of upstream intersection, or 15’ past the upstream transit stop (whichever is closest). Calculate this using a speed of 18 MPH for 25 MPH streets, and 22 MPH for 30 MPH streets.
- **Slack** – no longer used for local detection. (Blank out if previous timing card used slack.)

Recommended Parameters – Transit Detectors cont.

- **Detection zone (bus)** - Typically measured from the center of the previous intersection of travel to the center of the intersection, or (distance to upstream bus zone minus 15'). A 15' buffer is added when measuring from bus zones in case bus does not stop at the designated location.
- **ETA threshold (bus)** - Measure distance obtained from detection zone.

Remote Detection: Where used, add remote detectors with an estimate time of arrival (ETA) up to a third to half the cycle length (e.g., 30 to 45 sec if cycle length 90). Settings specifically for Remote Detectors.

- **Travel Time** – calculate similar to Local detectors, but add in dwell time for any intermediate transit stops plus 15 seconds for acceleration and deceleration.
- **Mode:** typically “raw check in”
- **Allowed late arrival to next detector** – this setting is typically set to 5, or can be calculated to be roughly 25% of the travel time to the next detector.
- **Time added if late** and **slack** – leave blank; we no longer use these settings.
- **IP addresses** for remote detectors should be included.

Adaptive Priority: D4's Adaptive Priority features no longer recommended and should be removed if included on a previous timing card change.

6. Preemption (Transit or Emergency Vehicle)

Preemption allows more aggressive measures to serve the transit or emergency vehicle signal phase as soon as possible and can include a variety of measures including deviation from normal phasing sequence (skipping or inserting phases), or holding the transit phase for extended periods. These modifications could require several subsequent signal cycles to step back into normal coordination.

Preempts are generally reserved for only the following circumstances:

- Cable car operations
- Railroad crossings and drawbridges
- Emergency vehicles with sirens active
- Transit operations requiring unique maneuvers (e.g. Market/Clayton)
- Actuated transit phases that cannot be accommodated in the traffic cycle

Preempts are different from TSP because they can:

- truncate (i.e. immediately end) Walk or Flashing Red Hand pedestrian phases.
- Insert special phases that can't be fit within the typical signal cycle
- Hold the desired phase for extended periods.

6.1. Preempt Settings

6.1.1. Truncating Pedestrian and Vehicle phases

In rare emergency situations, there is not enough time to serve the full Min Walk or Flashing Red Hand time prior to the arrival of the emergency vehicle or rail vehicle. In these cases, you can enable **Zero phase ped walk**, **Zero phase ped clear**, or **Zero phase green** to immediately terminate ped phases in preparation for a preempt event. For Emergency Vehicle Preemption, the ped walk phases are typically allowed to be aborted, but not other minimums. For more guidance and where these features can be used, refer to the CA MUTCD.

6.1.2. Track Clearance

Track clearance is typically used for heavy rail, but is also used in San Francisco for draw bridge preempts. Track clearance is used to clear the tracks (or bridge) of vehicles in advance of an event.

6.2. Getting Back in Step - Custom Transitioning

Locations that are anticipated to have frequent preempts, such as locations with revenue transit movements that are served by preempt, will frequently be thrown out of coordination and may not be able to recover within one cycle. The default for **maximum transition** and **minimum transitions phase splits** is to allow the controller to go “long way” by 25 to 33% of the cycle length, but only 6 seconds “short way”. This makes it more likely that a signal will need to go “long way” into step. If this happens frequently, it may increase delays for traffic and pedestrians on cross streets. This can be mitigated by entering more aggressive transition splits, particularly by allowing shorter minimum transition splits.

Recommended Parameters – Base Timing Settings (for Preemption)

- **Min Green** for phases with an LPI that may be truncated due to preempt should be set using the formula (min green = FRH – LPI). Combined with flagging “early walk to green” on the Preemption page, this will allow the green ball to be shortened only to the minimum required for the FRH.
- If a signal is expected to have frequent preempts, such as revenue transit phases that are served by preempt, there are settings on the Base Timing page that also should be entered. Setting the **Maximum and Minimum Transition phase splits (“Min Trans” and “Max Trans”** will allow a controller to get back in step more quickly.
 - For cycle lengths of 90 or less, set the splits so that the controller can get back into step in no more than three cycles.
 - For cycle lengths over 90, the controller can take up to four cycles.
 - To determine how much to extend the transition cycles, divide the cycle length by the number of cycles allowed, then add 1 or round up. For example, for a 90 second cycle, the max cycle lengths be $[90 \text{ sec cycle length} + 90 / 3 \text{ cycles} + 1] = 90 + 31 = 121$ seconds.
 - If the min timing and traffic volumes allow it, you can also shorten phases, e.g., allow a min trans cycle of 75 seconds, and a max trans cycle of 121 seconds.
 - When deciding where to add or remove time from the cycle, consider the demand on each movement and allocate the time accordingly.

Recommended Parameters – Transit, Rail, Bridge Preemption

- **Movements** – Describe the movement (including phase #) for the preempt (“Heavy Rail” or “Northbound Van Ness Bus Exit”)
- **Description** – Describe how the preempt should function to ensure that the settings are set correctly. This section is useful for a layperson to understand how the Preempt works.
- **Track Clearance 1 / Track Clearance 2** – Enter "V", "VP" or "P" to indicate which phases to clear during Track Clearance 1. Enter "-" if not using Track Clearance. Leave blank if there is no such phase at the intersection.
Phase Early Walk to Green – Enter "X" for vehicle phases with a corresponding early walk interval. For fixed-time vehicle phases with this parameter flagged, also adjust the minimum green to equal the FRH minus LPI time (see Recommended Parameters - Base Timing)
- **Zero phase ped walk / Zero phase ped clear / Zero Phase green** - Enter "X" to flag a phase that can be aborted, or leave blank to time normally.
- **Dwell** – indicate V/P/VP etc. for the phases that should dwell during the preempt phase.
- **Exit Phase** – if the preempt should always exit to a specific phase, indicate those phases. If exit next, write “NEXT”.
- **Exit Mode** – Select from these options:
 - **Normal** – Exits to the exit phases indicated on the line above.
 - **Next** – Exits to the phase that would have been next before the preempt was served.
 - **IS (In Step)** – Rarely used. Exits to whichever phase would be displayed if the signal had continued to time normally based on the offset and base timing. (e.g. if it is 48 seconds into the cycle and 3rd Street usually displays at that time, it will exit to 3rd St.)
- **Track Clearance** – enter the duration of each track clearance phase
- **Dwell (min time)** – enter the minimum duration of the dwell phases (e.g. “17”)
- **Max Dwell** – Enter maximum dwell time (in seconds). Leave blank if preemption should not be aborted (heavy rail). If the dwell phase is a fixed amount of time (e.g. there is no exit detector), the Max Dwell should be at least 1 second longer than the min dwell (e.g. “18”)
- **Checkout Limit** - Enter checkout limit time (in seconds). This is used if preemption detection is via momentary sensors, e.g., cable car check-in/check-out and LRV Vetag.
- **Change Phasenext** - allows the controller to modify the phase that follows current phase
 - Typically enter "Yes" (default setting), but note it may lead to red reverting if the preempt phase is clearing at the beginning of the preempt.
 - Enter "No" if not desired. Change
- **Outputs** – leave blank unless told otherwise by electricians. Typically used when communicating with track switch controllers and needing to send a signal from the signal cabinet to the track switch controller to confirm whether certain steps have occurred.
- **Detectors** - Describe if known, including type (momentary or presence) and location/distance from intersection.

San Francisco provides Emergency Vehicle preemption (EVP) to San Francisco Fire Department vehicles equipped with emitters. The parameters specific to EVP are below.

Recommended Parameters for Emergency Vehicle Preemption

- **Movements** – Describe the movement including phase # for the preempt (“Northbound Van Ness Emergency Vehicles”)
- **Description** – Describe how the preempt should function. (see general preempt instructions)
- **Track Clearance 1 / 2** – Not used for EVP
- **Phase Early Walk to Green** – Enter "X" for vehicle phases with a corresponding early walk interval. For fixed-time vehicle phases with this parameter flagged, also adjust the minimum green to equal the FRH minus LPI time (see Recommended Parameters - Base Timing)
- **Zero phase ped walk** –Typically flag all ped phases “X” for EVP to allow them to be aborted.
- **Zero phase ped clear / Zero phase green** - Not typically used for EVP.
- **Dwell** – typically use “V” for the direction that is being called. If there is a parallel phase, V can also be indicated for that phase (e.g. if Preempt is being provided for NB Phase 2, a V can also be entered for SB Phase 6).
- **Exit Phase** – use engineering judgment. Typically, the next phase after the vehicle phase served by emergency preempt.
- **Exit Mode** – typically “Normal”
- **Track Clearance (duration)** – Not used for EVP
- **Dwell (min time)** – typically 10 seconds for EVP
- **Preemption Max Override** – typically 120 for EVP.
- **Checkout Limit** – typically leave blank
- **Change Phasenext** – Typically “Y”
- **Outputs** – typically leave blank
- **Detectors** – describe where the GPS detector is placed.

7. APPENDIX - Advanced Features Rarely or No Longer Used

7.1. Travel Time Slack

The Travel Time Slack feature is intended to address the uncertainty of the transit travel time, but was found to reduce the effectiveness of TSP and thus is no longer recommended. However, the concept is explained here to explain the rationale against using this feature.

If a value was used for travel time slack, a green extension would not be performed if the green extension could not capture the expected arrival **plus the travel time slack**, so that if the transit vehicle is slightly later than expected transit travel time the controller would be able hold the green extension longer. Using travel time slack that accounted for the potential for some buses to be late resulted in fewer green extensions because the request for an extension had to be received earlier in the cycle.

Example:

- With no slack: Just before the traffic signal changes to yellow, an approaching transit vehicle is detected and is expected to arrive in 8 seconds. The controller is able extend the green for a maximum of 11 seconds and no value is entered for travel time slack, so the green extension will be granted for that vehicle. If a bus is able to pass through the intersection with a typical travel time of 8 seconds (or is slightly slower at 11 seconds), the bus will have zero signal delay. But if the vehicle misses the green extension, it would arrive during the yellow or red phase. It may have to wait for the entire red signal, rather than arriving a few seconds after the light turned red.
- With slack added: Assume the same scenario as above, but with the controller programmed with a Travel Time Slack of 5 seconds. When the bus is detected and estimated to be 8 seconds away, the controller would check if the potential green extension can account for the travel time + travel time slack, or $8 + 5 = 13$ seconds. Even though the estimated arrival is in 8 seconds, the controller will not provide the green extension because the controller can only extend up to 11 seconds, not the calculated 13 seconds. By accounting for the possibility of some buses arriving late, **every** bus that arrives at this point in the traffic cycle would have to wait through the red light, even though most buses would be able to get through on an up-to-11-second extension.

7.2. Estimated Delay

The estimated delay is typically used in situations with estimated arrival times higher than a third to half of the cycle length. Given that we have found it ineffective to prepare for transit arrivals that far in advance, this feature is no longer used.

When a remote detector is actuated and the travel timer starts counting down, the controller will pause the travel timer if it determines that the transit vehicle is stuck at a red light. D4 recognizes if there is a signal is between two remote detectors, if the signal is on the same IP network. If the light

turns red when the LRV has checked in but hasn't yet checked out, an estimated delay will be added to the travel time based on the signal timing at the stopped signal.

Example: Remote detector 2 has a travel time of 50 seconds. The transit vehicle passed remote detector 2 and triggered the travel time counter to start counting down from 50, and is now nearing the intersection prior to remote detector 1. However, the transit vehicle gets stuck at a red light before it actuates remote detector 1. The controller can estimate that the transit vehicle is stuck at a red, so it will pause the travel time for the entire red delay duration. The remote detector 2 travel time will resume when it receives its green and remote detector 1 should be actuated soon after.

7.3. Adaptive Priority

Adaptive priority can dynamically adjust the programmed travel times by "learning" from previous transit vehicle arrivals. In theory, adaptive priority modifies the estimated arrival time as traffic conditions fluctuate, e.g. increasing travel time during congested periods, then lowering travel time as congestion eases. However, in practice the D4 application of adaptive priority did not improve TSP effectiveness and in some cases had negative impacts (increasing the estimated travel time, but never lowering it back down). For this reason, enabling using D4's Adaptive Priority settings is no longer recommended and should be removed if it was included on a previous timing card change.

With Adaptive priority, the controller will determine whether to increase or decrease the travel times when a transit vehicle has a "max out" (transit vehicle checks out later than expected) or "gap out" (transit vehicle checks out earlier than expected) by considering the below settings.

- The *detector slack* in adaptive priority should not be mistaken with typical local and remote detector slack. This adaptive slack value is the amount of time for the transit vehicle to be late or early to the checkout detector before a "max-out" or a "gap-out" condition is recognized. The timer for the slack value will pause if the transit vehicle is stopped at a red signal. The timer will resume as soon as the transit signal goes green, therefore care must be taken to set this value high enough so that the transit vehicle can proceed from a stopped position at the stop bar to the checkout detector without recording a "gap out" condition.
- The *adaptive threshold* is the amount of "gap out" or "max out" occurrences there can be before a travel time adjustment is made.
- The *step* is the number of seconds that will be added or subtracted from the travel time after each adaptive threshold is met.
- The *max* is the maximum travel time that the controller can add to its base travel time. The base travel time is noted on each local and remote detector. The travel time is never adjusted below the base travel time that is programmed.