

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Vince Valdes: ITS Standards can facilitate the deployment of interoperable ITS systems, and make it easier to develop and deploy regionally integrated transportation systems. Transit standards have been developed by transit professionals like you at a national level to encourage competition and limit costs within our industry. However, these benefits can only be realized if you know how to write them into your specifications and test them. There are now a series of modules for public transportation providers that cover practical applications for promoting multi-modalism and interoperability in acquiring and testing standard-based ITS Transit systems.

Patrick Chan: This is Module 20, which is the Application of Arterial Management and Transit Signal Priority Standards.

My name is Patrick Chan, and I've been involved with the development of ITS standards since the year 2000: first as a public agency representative doing dynamic message signs before being involved with other standards. Prior to working with the public agency, I worked for eight years on the New York City traffic signal control system. I was a consultant that developed version 2 of NTCIP 1211: Objective Definitions for Signal Control Priority, which is one of the standards that we're going to talk about today. I've also been involved with several other ITS standards, including NTCIP 1202, the ITE AASHTO Traffic Management Data Dictionary, and SAE J2735: Dedicated Short Range Communications Message Set Dictionary, which is the data dictionary for connected vehicles.

Patrick Chan: This is actually the third module on arterial management and transit signal priority standards. The first two modules are Module 8, which is focused on understanding user needs, and Module 9, which is focused on specifying requirements. It's recommended that participants take these courses prior to this one, so that they're familiar with some of the key concepts in transit signal priority systems, even though we're going to review some of those key concepts in this module. But the learning objectives for these modules are: to be able to specify and test the transit signal priority implementation; describe how transit signal priority may be provided in a connected vehicle environment; explain the role of transit signal priority in integrated corridors; and finally, we're going to review some case studies where standards were used to provide transit signal priority.

Patrick Chan: For our first learning objective, we're going to talk about how to specify and test a transit signal priority implementation. We're going to first talk about identifying some potential issues with the standards, and talk about how to test the standards-based TSP implementation.

Patrick Chan: But before we start talking about NTCIP 1211, version 2—which will be the first standard we'll talk about—we just wanted to review some key concepts about transit signal priority. First, we need to define the components of a transit signal priority. Those key components are: the Priority Request Generator, which is the source of the request with signal priority; the Priority Request Server, where requests are collected and then processed; and finally, the Controller, which is the traffic signal controller that controls and operates the transit signal at a signalized intersection. The second key concept we want to introduce is the ITS standards define six different architectures for implementing TSP—Transit Signal Priority. Each architecture is defined by the physical location of each of the three key components that we just discussed. And implementation is not required to use one of the six architectures that are defined in the standards. In fact, they may use a different architecture, or they can support multiple architectures

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

that are defined. But just to point out that these architectures were originally defined to serve as the basis for developing these ITS standards.

Patrick Chan: We'll first start with NTCIP 1211, version 2. One of the first issues is that when generating a request for signal priority in NTCIP 1211, one of the things that must get sent along with the request is something called priority strategy. The priority strategy defines which approach the requesting vehicle will enter, and which approach the requesting vehicle plans to exit an intersection. This actually requires that the vehicle maintain a database of the priority strategies for each intersection that we would like to request priority from. So if it's the bus that's generating the priority request, then the bus maintains such a database with all of these priority strategies.

Patrick Chan: The second issue is that the requesting vehicle must know what vehicle type and vehicle class it is. The vehicle type and vehicle class is used to determine priority of a vehicle in case there are conflicting requests received at an intersection. For example, who gets priority if I get a request from an ambulance on one approach and the transit vehicle on the second approach? That's what the vehicle type and vehicle class is used for. The vehicle type defines what type of vehicle it is. Is a public safety vehicle? Is it an ambulance, a fire truck, or police vehicle? Or perhaps it's a transit vehicle, or perhaps it's a freight vehicle. The vehicle class, on the other hand, is a specific category of the vehicle type. For example, a transit vehicle: we may have different classes of transit vehicles. It could be a bus or rapid transit. It could be a light rail. It could be express service. Or it could be a local service. By establishing the vehicle class, we can establish, "Hey, an express vehicle gets priority over a local transit service," for example. The problem occurs when we have a regional implementation. By that, we mean that there are different jurisdictions, different agencies that operate the traffic signals, such as a local, a county, and a state transportation agency. And we also have different agencies that operate the vehicles that may request priority. For example, public safety, the fire department, the ambulances, different transit agencies. They all have to come to an agreement on which vehicle types will be supported, and which vehicle classes will be supported. We have to define them. So in a regional implementation, the regional agencies have to come together and agree. These are the vehicle types and this is what they mean; these are the vehicle classes, and this is what they are, and who gets priority. And this could be really an issue for transit agencies, for example, that operate in different traffic signal jurisdictions.

Patrick Chan: Those are two issues of NTCIP 1211.

Patrick Chan: Another transit signal priority standard is TCIP, or the Transit Communications Information Profiles. This is maintained by APTA—the American Public Transportation Association. It's used to focus on supporting business systems for transit systems. For example, it defines how data is exchanged between the different transit subsystems, such as a scheduling system, an operations system, and maintenance system. The issue, unfortunately, with TCIP is that it's just not widely deployed in the United States. TCIP is very complex, so it's very difficult for medium-sized agencies or small transit agencies to deploy and maintain TCIP. Determining which system architecture that could be implemented by a TSP implementation is probably the most important step for using the TSP standards. Overall, there are many different ways to implement transit signal priority, but by specifying which transit signal priority system architecture to implement and to support is probably the key step. Implementation can support just one system architecture or it can support multiple architectures, and it doesn't have to use one of the system architectures that are defined by standards. They could develop their own.

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Patrick Chan: The next couple of slides, we're going to talk about how to test the standards-based transit signal priority implementation. Why do we perform testing? Well, the easy answer to some people is that just because it's a payment milestone, but it's really more than that. We perform testing really to identify errors and bugs so that they can be corrected before the system is actually deployed. More technically, we perform testing to verify that this system was built correctly. Does it fulfill all the requirements that are in the project specification? So it answers the question: Was the system built right? Another reason why we perform testing is to validate that the right system was built. Someone can provide a system that's really nice—has all of the nice bells and whistles and looks cool—but we built the system, we wanted to procure a system because we had a transportation problem to solve, and if the system doesn't help us solve that problem, what good is the system? So one of the reasons why we do perform testing is to confirm was the right system built? Did we build the right system to solve the transportation problem that we were trying to address?

Patrick Chan: What we're showing here is something called a systems engineering V diagram. I'm sure many of you are familiar with it, but what it really represents is the lifecycle of a system—in our case, a transit signal priority system. It starts on the left side on top with when the idea for the system came about. It came from the regional ITS architecture, or it could have just been an idea that someone had for a project. Going down the left side, the next step really is to develop and determine what our user needs are. What is that we need the system to do? And that's documented in the concept of operations. Based on the concept of operations—the user needs—we develop system requirements. What do we need the system to do? Then we can develop a high-level design and a detailed design. These concepts were originally discussed in the first few modules on arterial management and transit signal priority.

Patrick Chan: The next step after that is we actually build the system. We go ahead and deploy the system, and then as we're deploying it, we actually start entering the testing phase, which is seen on the right hand of the V diagram. During this testing phase, we start asking: Can we verify and validate that the system that we wanted to build—was it built correctly, and did we get the right system? So for each step in development, there's some type of testing that we can perform along the way so we can verify and validate the system that has to be built.

Patrick Chan: Just to go into a little bit more detail. The first three are really a verification. It's an ongoing process. We verify that the requirements are fulfilled—that is that we built the right system. There's Unit/Device Testing where we focus on the specific component—whether it's a physical device, or maybe a specific interface. Then we may do a Subsystem Verification where we consider other environmental factors—such as the communications and how a component interacts with other components immediately around it. And finally, we do a System Verification and Deployment, where we consider the overall system and all the different components that make up a system and say: Does the system work properly as a whole, and does it work together to fulfill all our project specifications or the requirements in the project specifications?

Patrick Chan: Finally, we do Validation. It answers the question “I can operate the system and satisfy all of my stakeholder's user needs? Did I build the right system?” The system is considered validated when it's approved by all the key stakeholders and agencies, when all of the project requirements are fulfilled, and when corrective actions have been implemented for any anomalies that have been detected.

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Patrick Chan: What exactly are we testing though when we perform testing? Well, there's two things that we are really testing. One, we're testing for compliance with the procurement spec. Did we fulfill all the requirements that are in the procurement specification in our RFP? But we're also testing for performance of standards when we're talking about standards-based transit signal priority implementation. Does the system fulfill all the requirements that were selected for the system as specified in the standard? And does it also conform with the other specified requirements of the standards of references? For example, NTCIP 1211 specifies using a communications protocol, which is NTCIP 1103. When we perform performance testing, we're testing not only do we conform to NTCIP 1211—the signal control priority object definitions—but do we also conform with NTCIP 1103 as required?

Patrick Chan: But we're also testing that the data exchanges occur as defined by the standard. The standard does define what sequence of events are supposed to occur to fulfill a specific requirement. We're testing to say: Did we follow the sequence of events that's in the standard? If the data content that's being exchanged conforms with the standard, do we properly handle any error messages? So if there's a problem the standards may indicate what do we do and how do we handle these error messages. Do we conform with that standard? And we're checking that the correct structure of the data—the data elements—are also fulfilled.

Patrick Chan: So recalling the structure of NTCIP 1211, version 2: NTCIP 1211 has system engineering content, which means that it defines the user needs that are supported by the standards. Based on those user needs supported, it defines requirements. These are the requirements that have to be fulfilled to say we satisfied a specific user need or a feature. Some of them may be mandatory: You have to fulfill this requirement. Some of them may be optional: It's up to you—up to the implementation—whether that requirement should also be selected for your implementation. And then for each requirement that is supported by the standard, it defines a single design for each requirement, and that design will include there's the data sequence and these are the data elements that have to be supported. As a note, some standards do include test procedures to verify if the implementation fulfills the requirements and specific standards. Unfortunately, NTCIP 1211 is not one of those standards. It doesn't define any test procedures. However, NTCIP 1211, version 2 does provide two matrices which assist us to support the use of the standard and support testing in the standard. The first matrix is called the protocol requirements list. It's a table that lists all of the user needs features that are supported by the standard; whether that user need is mandatory to be supported to conform to the standard; and, for each user need, what are the requirements that have to be fulfilled to satisfy a user need.

Patrick Chan: What does this have to do with testing? The PRL, based on what user needs are selected by an implementer, supports testing by identifying which user needs have to be tested. For the example shown, we have a user need called Determine Priority Request Criteria, and this user need has three requirements that have to be fulfilled. Those requirements are Retrieve Priority Request Settings, Retrieve Reservice Period for a Vehicle Class, and Retrieve Priority Requests Time to Live Value. This user need allows an operator to determine how a priority request server is programmed to handle priority requests. What does it do, and what type of priority requests can it handle? The first requirement is to retrieve the setting: What are its settings? The second requirement is to retrieve the reservice period. So a reservice period is a time that must lapse after a priority request is serviced before another priority request is serviced. That means after a service provider priority request for someone, we want to wait a minimum period of time before we handle another priority request. Otherwise, the signal control will

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

always constantly be accepting priority requests, which can screw up coordination for the traffic signals and disrupt traffic flow. The last one—Retrieve Priority Request Time to Live—is what's it set up for? How long will we consider and support a specific priority request? Because if a vehicle sends a priority request, we may say the time to live is two minutes, so after two minutes, if I don't hear any updates or any other information from that vehicle for that priority request, I'm just going to start ignoring it. So a request isn't trying to be handled by a signal controller for a long period of time. Those are the user needs and requirements. These are the three requirements for this particular user need. The Conformance column says for the user need, determine priority request criteria. Is it mandatory? Does this user need have to be supported? Is it mandatory to conform to the standard? And then the next three mandatory says if this user need is selected, then these three requirements also have to be supported. And support is where an implementer can circle if is this required to be supported for my specific implementation? If it's optional, the support column will say "Yes/No," and then to specify this indicates whether that particular user need or requirement should be supported, or has to be supported, for your implementation. If it's yes—you want it to support that feature requirement—you circle "Yes." If you don't need it for your implementation you would circle "No."

Patrick Chan: And by the way, the student supplement has a more complete example of what we're going to show in these upcoming slides on how to use the PRL, and the second matrices—which we're showing now—which is the requirements traceability matrix. The second matrix is the Requirements Traceability Matrix indicates for each requirement the standard design that must be used to fulfill the requirement. The Dialog ID indicates the sequence of events that are expected to occur, while the Object identifies the object names to indicate the objects to be supported to fulfill the requirement. Note that the objects, we don't have a completed object on this for retrieve reservice period for a vehicle class Justin the interest of space. These two matrices—the Protocols Requirements List and the Requirements Traceability Matrix—are key and very helpful for the development of test plans and test documentation for NTCIP 1211. With the information in these tables, if test says know where requirements should be tested on the PRL because it indicates what user needs and requirements are to be tested, and then Requirement Traceability matrix indicates that for each of the requirements, what is that we must test to verify if a requirement is fulfilled. So, for this first requirement, we'll indicate you have to fulfill the requirement retrieve priority request setting, you have to use dialog 4.3.1.1, and you have to support the object ID 5.1.2.7, which is prsProgramData.

Patrick Chan: Based on these two matrices, an implementer can now create something called a Requirements Test Case Traceability Table. So what this is is a table that an implementer can create which lists out all the requirements from the standard to be tested. In our example, we're showing these are two requirements that we want to test: Requirement 3.5.1.3.1: Retrieve Priority Request Settings, and 3.5.1.3.2: Retrieve Reservice Period for a Vehicle Class. Then below that for each requirement it defines the test case that has to be performed and passed to verify that we fulfilled that requirement. This table is helpful in that it verifies that at least one test case is identified for each requirement that we want to test, and that we know which test cases have to be passed so that we can verify that the requirement is fulfilled. Note for the second requirement—Retrieve Reservice Period for a Vehicle Class—there are two test cases. An implementation must pass both of these test cases before we consider the requirement is fulfilled.

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Patrick Chan: Multiple test cases may be needed to completely test the requirement. Each test case may have a different set of values; each test case may test different conditions. Each test case should confirm that the interface performs at the same sequence of data exchanges or events as defined in the standards—that's the dialog—and uses the data concepts—the data objects—that are indicated in the RTM. The student supplement contains examples of more complete Requirements Test Case Traceability Tables, and also has an example of what a test case may look like.

Patrick Chan: We've reached our first activity. The purpose of the activity is reinforce what we've learned so far in the learning objectives.

Patrick Chan: Our activity is the question: "Which of the following is not a reason to perform testing?" Your answer choices are: A) To identify bugs or errors so they can be corrected; B) To verify the system fulfills the requirements of the specification; C) To validate the right system was built; and D) To check a box that we did it.

Patrick Chan: The correct answer is actually D: To check a box that we did it. So we perform testing not to satisfy chronological lists, but really to make sure that the system that we wanted—that we tried to procure—is provided. We do use testing to identify bugs or errors so that it can be corrected before we deploy it; to verify the system fulfills the requirements of the specification; and just to validate that we built the right system to solve the transportation problem we were trying to solve.

Patrick Chan: For the next couple of slides, we're going to describe how transit signal priority may be provided in the connected vehicle environment.

Patrick Chan: First, we're going to discuss briefly what a connected vehicle environment is. This provides a context for a connected vehicle environment. In the connected vehicle environment, a vehicle can broadcast information wirelessly about itself. Information that may be broadcast is—this is where I am right now, and these are the values that's on the sensors on my vehicle. The connected vehicle environment, these vehicles can broadcast this information to other vehicles around it, or maybe to the infrastructure, such as roadside equipment or to traffic management incentives, or even devices such as smartphones. In return, a vehicle can receive information from other vehicles or from the infrastructure. By getting this information, we can reduce the likelihood of incidents as vehicles now know where each other are. I know I'm driving down the roadway or the freeway or the arterial and I can tell wirelessly, through this connected vehicle environment, that there is a vehicle in my blind spot on the left that wants to make a right turn in front of me because I'm at a signalized intersection. It can also improve mobility—that is reduce delays—as we have more information about the traffic, the transportation environment around us. It doesn't have to be a vehicle. It could be a smartphone that's on a pedestrian or a bicyclist that's also sharing and providing this information.

Patrick Chan: There's a little diagram that kind of helps you visualize what a connected vehicle environment is. The vehicles are broadcasting information about themselves. It might be broadcasting this is where I am. This is the direction that I'm moving in. This is my speed. Whether my turn signals are on to indicate I want to make a left turn or a right turn. What the length or my width of vehicle is so that you know am I a 55-foot transit vehicle, or am I just simply a 25-foot transit vehicle? In return, the infrastructure can provide my vehicle with information also. It can provide me information about the signal

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

phase and timing: There's a traffic signal here and it's about to turn green for the northbound direction, and it's about to turn yellow in the eastbound direction in two seconds. It can provide information about speed limits on the roadway, or the estimated time of arrival for transit vehicles. The point is that due to this connected vehicle environment, the service transportation infrastructure can provide information about the vehicles and the roadway. And all of this can be used to improve safety to increase the mobility and to decrease the environmental impacts.

Patrick Chan: So what's driving all of this? What's creating this connected vehicle environment? In August 2014 NHTSA—the National Highway Traffic Safety Administration—released an Advanced Notice of Proposed Rulemaking and a supporting research report. What the proposed rulemaking is that in the future, all new vehicles will have to be equipped with vehicle-to-vehicle communications capability—all new light vehicles, excuse me. That way passenger cars can broadcast information about itself, and also receive information from other light vehicles—passenger vehicles. The research report that was also released concluded that a fully mature V2V system—vehicle to vehicle system—alone can potentially address 79 percent of all vehicle crashes that occur in the United States, and that a vehicle-to-vehicle and a vehicle-to-infrastructure environment concurrently can potentially address 81 percent of all vehicle crash types. The rulemaking on the vehicle-to-vehicle capability is to be required on light vehicles is expected to be released sometime in 2016. Now this does not affect transit vehicles specifically, but that doesn't preclude transit vehicles from being equipped to provide, or to have, V2V communication and support.

Patrick Chan: This is important because the vehicle-to-vehicle communications, when required, will open the gates for lots of different applications or opportunities for other V2X communications. Once it's available, infrastructure owner-operators can take advantage of this vehicle data to receive information from the vehicle—so now we're more aware of what's happening on the transportation infrastructure—but also to provide information to the vehicles in the interest of trying to improve the roadway safety overall. Improve mobility so they'll reduce congestion, and to improve the environment so we can increase fuel efficiency and reduce recurring congestion.

Patrick Chan: What information is exchanged for transit signal priority in a connected vehicle environment? Well, that's really defined in one of the interconnected vehicle standards called SAE J2735, which is Dedicated Short Range Communications (DSRC) Message Set Dictionary. This is the ITS standard that defines the dictionary, the words, the messages, the data elements for connected vehicles. Some of the messages that it supports is MAP data message, which provides information about this is what the roadway looks like, this is what the intersection looks like. Here are the lanes—and this lane is for transit vehicles only, or this lane contains a transit stop. The Signal Phase and Timing: this provides information about the signal timing at an intersection. Which direction it gets green. Which is it: a left turn that's providing green, or pedestrians are allowed to cross at this intersection, and which way? It also has a message for signal requests. This is Signal Request Message, where we could request signal priority, and Signal Status, which is the status of the signal priority requested. It's SAE J2735 that defines the messages for connected vehicles. So in this graphic, we're just depicting a quick transit vehicle approaching an intersection, and its broadcasting and receiving its position. And on the left, we also have a DSRC roadside equipment which is broadcasting information from the infrastructure to vehicles.

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Patrick Chan: Since we're focusing on transit signal priorities, the Signal Request Message is of most interest to us. It's broadcasted by a vehicle to the infrastructure. And within the Signal Request Message, we're asking for transit signal priority, which includes priority and preemptive treatment from one or more traffic signals. The SRM is the request message for signal priority and signal preemption. The message allows the vehicle to identify itself, specify—as a request and priority of pre-empts—signal priority or signal preemption. It includes information about the vehicle speed, heading, and location.

Patrick Chan: Some of the mandatory elements in the Signal Request Message—that means that these are elements that have to be included and this Signal Request Message—is the Requestor Identifier of the vehicle, for example. A Request Identifier. This is my third request, so a fourth request on my run, for example. The Request Type. Is this a new signal priority request? Or are we updating an existing one that we've already sent out, because we had to stop for a vehicle so we have to change our estimated time of arrival? Or is it a cancel? Well, I've already gone through the intersection and so please go ahead and cancel my transit signal priority. Another mandatory element is what lane or approach that I'm approaching the intersection. So I'm in lane five, which is going in the northbound direction, for example. Some optional elements that can be provided as part of the Signal Request Message is the Estimated Time of Arrival. This is what time I expect to arrive at the intersection. Estimated Duration. This is how long I think before I arrive at the intersection and how long it's going to take me to go to the intersection. Some information about the requester. Is it currently in the middle of a transit route, for example? Is it a transit vehicle? And because I'm running behind schedule, please increase the priority level because I'm a full bus carrying passengers. In addition, there's other optional elements that are supported in the SAE J2735 in the Signal Request Message specifically for transit. Is the relative passenger accuracy? Is it an empty bus? Is it a full bus? Is it currently on schedule for the schedule adherence? Is it currently stopped or loading passengers with the door open? These are information that can be used and provided by the transit vehicle to the signal controller that can help it determine what type of signal priority to provide.

Patrick Chan: Another message is the Signal Status Message, which is broadcasted by the infrastructure to everyone—to all the vehicles. Once it receives signal priority requests, the signal status tells the vehicles, "This is the status of the priority request." It's broadcasted for the entire intersection, and it sends out the current status of any priority requests or preemption requests that the signalized intersection has received. This message includes the outcome of the pending request—we're going to grant it or not—and it contains the identifier of the intersection and the status of each request for a specific lane or for a specific approach, depending on how the signalized intersection is configured.

Patrick Chan: So just to give an example, this is the example scenario or use case of how transit signal priority would work in a connected vehicle environment. A transit vehicle is approaching a signalized intersection and it enters the range—because remember, it's wireless—enters the range of the RSE and the signal controller. The transit vehicle then wirelessly broadcasts the signal request message with its estimated time of arrival and the identifier of the lane that it wants to enter and exit out of the intersection. The signal controller then receives and processes the signal request message.

Patrick Chan: Next, the signal controller then provides the RSE—the roadside equipment, which is the radio—with the signal status message so that the roadside equipment can broadcast the signal status message with the status of all the signal requests that the signal controller has received. It will broadcast

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

that signal status message. The transit vehicle will receive that signal status message, and then travel to the signalized intersection when the service is provided.

Patrick Chan: We've reached our next activity.

Patrick Chan: And the question is "Which ITS standard defines the messages and data elements for a connected vehicle environment?" Is it NTCIP 1211, version 2? Is it SAE J2735?—that's B; C is TCIP; and D is NTCIP 1103. So which ITS standard defines the messages and data elements for the connected vehicle environment?

Patrick Chan: The correct answer is actually B: SAE J2735. SAE J2735 was developed specifically to support a connected vehicle environment. It contains the messages, the data frames, the data elements for the connected vehicle environment. NTCIP 1211 supports transit signal priority, but not necessarily for a connected vehicle environment. TCIP supports the transit business, but not specifically for a connected vehicle environment, and NTCIP 1103 defines protocols for managing transportation field devices.

Patrick Chan: Our third learning objective is to explain the role of transit signal priority in integrated corridors. For the next couple of slides, we're going to discuss the impact of transit signal priority on integrated corridors.

Patrick Chan: First, let's review what is an Integrated Corridor Management, or ICM? Integrated corridor is where we take a look at the existing transportation infrastructure along transportation corridors and try to optimize it. Meaning, try to use the full capacity regardless of the mode of transportation—whether it's the freeway, the arterial, and the transit resources that we have along a corridor—so that we can try and fully use the capacity of that transportation corridor. One of the benefits of doing Integrated Corridor Management is that it enables travelers to make informed travel decisions and dynamically shift modes for transportation, even if they're in the middle of a trip. The purpose of the integrated corridor is just to reduce travel overall: transportation, reduce travel time, delays, fuel consumption, emissions, and incidents.

Patrick Chan: So how can transit signal priority contribute to an integrated corridor? Recall that one of the benefits of transit signal priority is travel time reliability, which is really important. Travelers that select transit as an alternative route, they were more likely to return and use transit again if they trust the predicted travel time to the final destination. So we believe that it's reliable. If they perceive if this is the estimated time to reach our destination and it's reliable, then passengers are more likely to use transit again. Transit signal priority can also decrease travel time on transit routes along the corridor by providing preferential treatment to transit vehicles—particularly if the transit vehicles are full. So by support, transit signal priority can also increase the capacity of the transit route—and throughput of the transit route—since the run time's more reliable and thus can be shorter. We can also use transit signal priority to enforce changes to transit schedules by providing or removing signal priority.

Patrick Chan: Some examples where transit signal priority has contributed to integrated corridors are in Dallas, Texas, along US 75; and also Minneapolis, Minnesota.

Patrick Chan: We've reached another activity.

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

Patrick Chan: “Which of the following is not a benefit of using Transit Signal Priority in Integrated Corridor Management?” The answer choices are: A) Decrease travel times; B) Improve travel time reliability; C) Improve the quality of transit data collected; and D) Improve throughput and use of transit capacity. Again, which of the following is not a benefit of using Transit Signal Priority in an integrated corridor?

Patrick Chan: The answer is C) Improve the quality of transit data collected. The quality of transit data is unrelated to the transit signal priority, but it has been demonstrated that transit signal priority can decrease travel times along a corridor, and it can improve travel time reliability. That's probably the most significant benefit of transit signal priority, especially to transit customers. To improve throughput in the use of transit capacity when we give it a priority we can contribute, because of the improved liability and decreased travel time we can increase the throughput and the use of the available transit capacity.

Patrick Chan: For the next couple of slides, what we plan to do is just review some use cases where the transit signal priority standards were used to provide transit signal priority. We're going to go through four different use cases, and each of these use cases has something different. How they implemented transit signal priority was slightly different—whether they used different transit signal priority standards or how they approached the implementation of the transit signal priority based on existing conditions.

Patrick Chan: The first case study example we'll go through is for King County Metro, which operates the transit system in the Seattle Metropolitan region.

Patrick Chan: What we're showing here is a map of the region and the locations of the six bus rapid transit corridors that currently exist in the Seattle region. There's also two non-bus rapid transit corridors, with more being planned. In King County Metro, they have approximately 200 signalized intersections which are currently supporting transit signal priority, and there were 13 local partner jurisdictions that were involved with supporting the transit signal priority implementation for King County Metro.

Patrick Chan: One of the cases to assess King County Metro's implementation was the design of its system architecture. It can support any of the five standard system architectures that were defined in TCIP. In their implementation, the bus notes its location and the location and the priority intersection approaches—how it's approaching the signalized intersection that supports signal priority. It initiates the request based on its current location. As a note, the conditions of its transit signal priority, if needed, can be easily changed, either based on the route or based on the time of day on each vehicle. Their signal priority implementation can also accommodate complex strategies, such as check in, check out, and near site stops. Their specific implementation has a separate device called the Transit Priority Request Generator that transit vehicles communicate with directly in a static device in the signal cabinet that generates the transit signal priority.

Patrick Chan: This is a picture of their communications network that was used for King County Metro. It's a unified ITS communications network. In addition to supporting transit signal priority, it supports other transit business systems, such as the fare payment, passenger counting, and providing real time information signs to passengers at the bus stops and bus shelters. It's a wireless communications network using a restricted frequency, meaning that no one else is on that particular frequency. There's no interference from other potential users; it's strictly used for King County. Where possible, they also use

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

common network components. They use standards where they're possible, so they were able to use the common standardized network components, which made it easier in terms of upgrading, maintenance, and especially with costs, since they can purchase off the shelf equipment. They also developed standard design and installation options for each site, so by having these standard drawings, every time they upgraded or added a new signalized intersection location, they had standard drawings that they could use to build out their infrastructure.

Patrick Chan: In terms of ITS standards, King County Metro supported the full TCIP dataset for transit signal priority. The request messages that were generated included the 25 defined fields in TCIP, plus 10 additional user-defined fields. The strategies and conditions are loaded in the priority request generator and the priority request server, and timing plan specific to support the transit signal priority were loaded onto the signal controllers. Each generator, every request message that was transmitted was logged and stored, and the same thing on the controller. Every time the controller received a request, they logged what phases, what phase was used to provide a signal priority, what the priority request generator did. It logged the priority type. It noted what type of transit signal priority was supported.

Patrick Chan: Some of the lessons learned from King County Metro. They found that the system engineering process was very helpful, especially with writing specifications and for testing. So King County Metro actually followed the V diagram. They developed signal concept of operations. They developed a requirements document, and that was very helpful for creating the project specifications and for when they actually had performed the testing. They discovered that transit signal priority algorithms varied between vendors, so each vendor—based on how the signal controller worked—had its own algorithm for providing transit signal priority. They took advantage of standards wherever available, especially with their communications network. They used IP/Ethernet where possible for their communications network, so it made the design implementation and operations and maintenance more cost effective. And because they had their own communications network—it was a high bandwidth that wasn't being shared with anyone else—so there was no limit to the size or frequency of the data, which was very helpful. There was no real-world constraints on their network in terms of supporting the transit signal priority.

Patrick Chan: The next case study we're going to talk about is New York City. In New York City, they maintain and operate nearly 12,400 signalized intersections, nearly all of them which are under computer control. MTA is the main operator of the transit buses in New York City. They have a fleet of approximately 6,000 buses—approximately 5,000 buses may be on the street at a time during the peak periods. All their buses are equipped with GPS and wireless communications with their transit management center, which is a big benefit they took advantage of, and we'll discuss that in the next slide. New York City had gone through some previously tested transit signal priority in two different corridors with successful results, resulting in decreases of 15 to 23 percent in travel time for their transit vehicles along those routes. They were convinced about the benefits of transit signal priority.

Patrick Chan: When designing their system and communications architecture and network, the New York City DOT had a concern about the capital costs. We were talking about 6,000 buses and 12,400 intersections. To add equipment for each of those buses and for each of those traffic signal controllers is going to be really costly, so it was cost prohibitive. Their solution was to equip the transit vehicle with a MAP database that involved no new hardware. By providing a MAP database for each transit vehicle, the

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

transit vehicle knows which intersection that vehicle is approaching. It knows where to transmit a request saying, "I'm approaching this intersection and I need transit signal priority." It knows where it is and where it where transit signal priority is supported. It knows where to transmit a clear—it means, "I've gone through an intersection, so go ahead and clear my transit signal priority request." And thus, it also knows which priority strategy to request. The priority requests are generated—transmitted wirelessly—from the transit vehicle to the transit management center. That request is then forwarded to the traffic management center, which then evaluates each request based on the business rules agreed upon by New York City DOT and MTA. If approved—meaning if we decide if the server agrees to provide transit signal priority—then it forwards the request to the traffic signal controller or to the traffic signal control system. They took advantage of their existing communications network to implement transit signal priority.

Patrick Chan: In terms of ITS standards, New York City DOT was already using NTCIP for communications between the traffic management center and the signal controller. There's a separate NTCIP doc standard called NTCIP 1202 that New York City DOT was already using to support the operations and management of their signal controllers. They adopted NTCIP 1211 objects with extensions. By extensions, I mean there was some functionality that New York City wanted—or some information that they wanted—such as vehicle speed, the vehicle location, a route identifier, the identified intersection that signal priority was being requested for but wasn't supported by the existing standards, such as NTCIP 1211. These were objects that New York City created to support their implementation, because it wasn't supported by the existing standard. As a note, New York City DOT is not conformant to NTCIP 1211, meaning it doesn't support all the user needs or requirements or data elements that are required by NTCIP 1211 to conform to the standards, but they do use the NTCIP 1211 objects.

Patrick Chan: Some of the lessons learned from New York City is they took advantage of their existing communication infrastructure. They already had communications between each vehicle bus, each transit bus, and the transit management center. They already had communications between the traffic management center and the signal controller. They just built out a communication network between the traffic management center and the transit management center just to decrease their costs. They also discovered that there was an issue with NTCIP 1211, version 1. This is version 1—we're now in version 2—in that there was a communication network latency, and that had a big impact on the system. The way version 1 worked is that the requester—let's say the transit bus—would say, "I'm going to arrive at the intersection in five seconds," so let's say five seconds after midnight. Because of the communications network—they were using cellular data between the transit bus and the transit vehicle and the transit management center—sometimes it would take a couple of seconds for the transit management center to receive the message. For example, if there was a communication network latency of five seconds, by the time the traffic signal controller received the message—like "I'm arriving in five seconds"—the transit vehicle was already there because of the network latencies. This was an issue that they had to deal with. But this ended up being addressed in NTCIP 1211, version 2. That issue doesn't hopefully no longer exist. A second issue that New York City learned was that the clock source can have an effect on the implementation. In the transit vehicles, the clocks on the transit vehicles was based on GPS time. So if GPS time said "12:00 midnight," that's the time that the transit vehicle thought it was. However, the traffic signal controller was on a different time source. It was actually on the Eastern electrical grid. It was based on the electrical power being provided by the electrical utilities. The time sources for the traffic signal controller and the transit vehicles were different. When the times were the same it was an issue. It was

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

possible that they may have different times but the source of their clocks were different. That's an implementation issue that you should be aware of.

Patrick Chan: The next case study example we'll review is for the Chicago Metropolitan area. They have a project called the Regional Transit Signal Priority Implementation Program. The main members of this project includes the RTA—the Regional Transit Authority in Chicago—the Chicago Transit Authority, PACE—which is a suburban transit provider—Illinois DOT, Chicago DOT, and several county DOTs and other municipalities in the Chicago Metropolitan region. The goal of this project was to develop and implement a regional transit signal priority system for Metropolitan Chicago. They've dedicated \$40 million dollars—which involves 100 miles of roadway, 400 signalized intersections, 13 arterial corridors and four counties for this transit signal priority project.

Patrick Chan: They had previously done some demonstrations of transit signal priority for the Chicago Transit Authority and PACE, and they were happy with the results because it improved schedule adherence—it reduced travel time up to 50 percent. They did discover that it was difficult to evaluate the performance of the TSP system because the separate vehicle logs and the control logs—they had separate logs—and it was just very difficult for them to correlate the different logs to see what did happen. Which vehicle was requesting this transit signal priority, etc. It was just difficult to correlate. Because of the regional nature of the proposed transit signal priority system, the region wanted to use open standards—so they weren't tied to a specific transit signal priority vendor, and to simplify operations and maintenance for the different agencies that were involved—and allowed PACE and the CTA to request transit signal priority from a single device within the signal cabinet, and also finally to provide centralized monitoring of the transit signal priority activities from their transit management centers. The benefit of doing the open standards was there was interoperability between the two transit agencies and various DOTs that maintained the traffic signal controllers, and to provide flexibility and fleet deployment for the transit agencies so that a vehicle would only work with one specific corridor for transit signal priority.

Patrick Chan: This is a quick overview of the signal architecture that was designed and implemented or planned for Chicago. It's primarily vehicle-to-infrastructure, but the PRG—the priority request generator—is located in the transit vehicle. It may be a separate physical device, or it might be part of an existing AVL system. The priority request server functionality is also being incorporated in the AVL systems, but also might be located within the controllers themselves. There might be a separate priority request server device, or it might be within the traffic signal controller device itself, as Chicago DOT right now is implementing and deploying a new advance traffic controller for their traffic signal control system. There might be separate PRS devices in the controller cabinet for legacy controllers—for the older controllers. Their system architecture allows traffic management centers to monitor transit signal priority operations via the signal controller. The architecture also is designed to support center-to-center communication so that the transit management centers can monitor the transit buses' locations and send requests directly to the traffic management center, which can then update the signal timing plan as necessary.

Patrick Chan: In terms of the ITS standards, when they first started the project, they spoke with the various stakeholders about their specific user needs, and based on that, they developed a set of requirements. And based on those requirements, they ended up developing a regional message set between the vehicle and the intersection based on NTCIP 1211. They took advantage of some information that was in SAE J2735 standard to fulfill some of the requirements that weren't supported by

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

NTCIP, such as the vehicle location. Once they developed a proposed message set, they sent that message set to vendors to solicit their comments via Request For Information. Since NTCIP was even more easily supported by vendors at the time—because SAE J2735 wasn't mature yet—they developed their message set based on NTCIP objects. Following the system engineering process and what's been recommended for extensions for NTCIP, they also updated the dialog definition so that it's consistent with NTCIP 1211, but includes their extensions. They also developed test tools to verify the correct usage of the data objects, so if they have a common test tool, then all the vendors can test again to verify that their implementation does indeed conform with their regional message set. Chicago also took the extra step of obtaining a proprietary node for their object extensions. They actually have their own NTCIP node where they can define and maintain the regional message set for this project.

Patrick Chan: This is a list of the object extensions that the Chicago region created to support their implementation. It includes an object for the vehicle ID for defining which phase they would like to request for the transit signal priority; the location of the vehicle; the agency ID; whether a vehicle is late, behind schedule or not; the route ID; the run ID; and the relative occupancy of the transit vehicle that's requesting transit signal priority.

Patrick Chan: Some of the lessons learned from the Chicago implementation. Developing a flexible system architecture was very important. It supports the existing infrastructure while supporting future upgrades, capabilities, and expansion. In some areas where the transit signal priority is being implemented, the agencies are still using low bandwidth telephone lines for communications between the management center and the controller cabinets in the field. They took advantage of that opportunity to upgrade the communications to those signalized intersections to higher capacity communication lines. There's lots of agencies that were involved in the project, but the cooperation has been great. Because of the flexible system architecture, many of the traffic signal controllers are dated, but they have been testing some of the new advanced traffic controllers with the PRS functionality built into it. The architecture easily supports these new advanced traffic controllers while also supporting some of the legacy traffic signal controllers.

Patrick Chan: The last case study we're going to talk about is something called MMITSS, or the Multi-Modal Intelligent Traffic Signal System. The goal of this particular project—and it's a demonstration project—was to develop a comprehensive traffic signal system that services multiple modes of transportation in a connected vehicle environment. This is a demonstration site for connected vehicles for traffic signal controllers. Now, this demonstration project wasn't focused purely on transit signal priority. It also supported pedestrian priority, emergency vehicle priority, and actually freight vehicle priority—so it was a full demonstration—but the focus of our slides will be just the transit signal priority aspects. This demonstration project was performed in two different locations. One was in Arizona at Anthem, Arizona—which is just outside Phoenix, Arizona; and Palo Alto, California.

Patrick Chan: This is the system architecture for the MMITSS project, as it's known. The transit signal priority includes the equipped transit vehicles generating signal request messages out over the air. It's received by the roadside equipment, which includes a radio and a processor called the MMITSS roadside processor. In addition to receiving the signal request messages, it's the RSE radio that broadcasts MAP data, information about the roadside geometry, and also SPAT information—the information about the signal phase and timing at the intersection. The RSE receives these priority requests from the transit

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

vehicles and signal timing information from the controller, and then based on that, it tries to calculate an optimal signal timing schedule to the signal controller. Once the signal controller receives it, that information determines how to allocate green time within the optimal schedule. It will provide that information back to the MMITSS processor and the RSV radio, at which point the RSE radio will transmit a signal status message back to the vehicles at the intersection—such as the transit vehicle. Once the equipped vehicle—the transit vehicle—departs the intersection, it will send back a cancel signal request message back to the RSE in the MMITSS, saying, "I've gone through the intersection, so go ahead and cancel my transit signal priority request."

Patrick Chan: The MMITSS project actually used an older version of SAE J2735—a 2009 version—as part of its field tests, and they used the signal request message and the signal status message for the priority request. The MAP message—the roadway geometry—and the SPAT message—the signal for timing information message—were also tested as part of this demonstration project. As a note, for this project they did modify the signal status message to serve as an acknowledgement of the receipt of a signal request message, because without this modification, there was no method by which a transit bus can verify, "This transit, the signalized intersection has received my priority request." With this modification it knows, "The signalized intersection has received my request and it's doing whatever." The demonstration project also used NTCIP 1202 to exchange information between the roadside equipment and the controller. And it also used NTCIP 1211 object definitions to provide information about the vehicle class and vehicle type.

Patrick Chan: Some of the lessons learned from this MMITSS project. Independent analysis showed that the MMITSS applications actually worked pretty well—it was pretty effective. It improved the vehicle travel time and the travel time reliability. It reduced delay for equipped vehicles at the test location. Specifically, for equipped transit vehicles, it reduced delay by 8.2 percent. Simulations also confirmed the evaluation that transit signal priority effectively saved travel time for the transit vehicles on the corridor, but did occasionally increase the overall system-wide delay due to the release of reduced green times on the side streets.

Patrick Chan: So we've reached our last activity.

Patrick Chan: The question is: How can ITS standards be used in transit signal priority implementations? Your answer choices are: A) Extensions to an ITS standard can be used to satisfy a need not supported by the ITS standards; B) NTCIP 1211 v02, TCIP, and SAE J2735 must be used in TSP implementation to conform to the TSP standards; C) All messages and object defined in the standard must be used to conform to the standard; and D) An implementation is allowed to support only one of the system architectures defined in the standard. So again, how can ITS standards be used in TSP implementations?

Patrick Chan: The correct answer is actually A, extensions to an ITS standard can be used to satisfy a need not supported by the ITS standards. The ITS standards allows an implementation to define an extension if the user need is not already supported by the standard. If the user need is supported by the standard, it is expected to conform to that standard that you will implement the standard the way the standard defines it. B is wrong because not all three standards have to be used. Depending on your implementation, you may select only one of the standards. All the messages defined in the standard—it doesn't have to be supported. You only support those parts of the standard that are applicable to the

Module 20-- Application of Arterial Management/ Transit Signal Priority Standards

features that you're trying to provide in your implementation. And an implementation doesn't have to support one of the system architectures defined in the standard. You can define your own architecture or you can support multiple, if not all, of the architectures defined by the standards.

Patrick Chan: So just to summarize what we have learned in this module. First, we've identified some of the known issues with the existing TSP standards, but we've also identified how to identify requirements that are to be tested, then how to test standards-based TSP implementation. We defined what a connected vehicle environment is and what information is exchanged with regards to TSP in a connected vehicle environment. We've talked about and explained the impact of transit signal priority in an integrated corridor. And we've taken away knowledge of some of the case studies that we reviewed in this module on how to correctly use the standards to execute transit signal priority.

Patrick Chan: We want to thank you for participating in this module. There is a feedback link below so that you can provide us with your thoughts and comments. I just also wanted to remind you to please take a look at the student supplement. There's additional references in the student supplement. There's references about connected vehicle environment. There's links to the first two modules related to arterial management and transit signal priority. There's modules related to testing. There's modules related to connected vehicle environment. And if you're interested in more information about the case studies, there are references and links provided also in the student supplement. So again, thank you very much for completing this module.