

Simulating the effects of PTC systems on railroad operations

Ted Jones and Ian Martin
SYSTRA Consulting, Inc.

1. INTRODUCTION

Positive Train Control (PTC) systems ensure train safety by enforcing observance of signal aspects and maximum authorized speed, temporary speed restrictions and stops for interlocking signals displaying a stop aspect. To achieve this level of safety through an automated system, certain worst case assumptions must be made regarding train performance and track conditions that are not known to the system. There is also error that is inherent in train location information available to the PTC system which must be accounted for. For these reasons, trains will perform differently in the presence of a PTC system than they would if these operating conditions were left to be enforced by the engineer through adherence to rules and instructions.

These differences in train performance will have an impact on train operations in the future under various PTC implementations. There are also design and implementation decisions that can affect and potentially alleviate some of these performance impacts. Until these various PTC implementations are in operation the only way to truly know the extent of these impacts and the potential benefit of various design decisions is through simulation.

This paper will discuss the simulation of train operations under a system without PTC installed and under the same system with PTC installed. Specifically this paper will focus on the Advanced Civil Speed Enforcement System (ACSES) which has been installed on the North East Corridor and is currently in use by Amtrak.

2. PTC MANDATE

On Friday September 12, 2008 at Chatsworth, California where a Metrolink train collided head-on with a Union Pacific Freight train causing the death of twenty-five (25) people and injuring a further 130 people, In response to this collision the Rail Safety Improvement Act, 2008 was enacted.

This Act, in seven titles, focused on Rail Safety improvements. Title One contained ten sections relating to Administration, strategy, risk reduction, PTC, technology grants, statutory mandates and recommendations, rulemaking process, hours-of-service reforms, risk analysis and pilot projects. Section 104 defined the timescales for all Class One railroads, to include intercity and commuter railroads, to install a PTC system prior to

December 31, 2015. In addition to the Class One Railroads any line that carries hazardous material classified as poison or toxic-by-inhalation will be subject to the same requirements.

Section 104 of the Rail Safety Improvement Act 2008 describes PTC as:-

“POSITIVE TRAIN CONTROL SYSTEM.—the term ‘positive train control system’ means a system designed to prevent train-to-train collisions, over-speed derailments, incursions into established work zone limits, and the movement of a train through a switch left in the wrong position.”

PTC is further described as ***“An integrated command, control and communication system that enforces movement authority if the system detects that a train will exceed its designated limits”***

PTC is technology specifically designed to automatically slow or stop a train before certain types of accidents occur. Its purpose is to ensure safety by preventing trains from exceeding the maximum permissible speed; ensuring trains slow down for permanent and temporary speed restrictions and forcing stops at interlocking signals displaying a stop aspect. In this way, PTC is intended to prevent train-on-train collisions, derailment caused by excessive speed and unauthorized access to worksites where repairs are being carried out.

The concept for PTC relates to three main components for vital and non-vital, train control and signaling.

Vital components assume a safe state to a very high probability for any and all independent failures including hardware and software and self-revealing failures. Non-vital is defined as a system with components that are not considered vital.

The signaling component conveys information to the engineer which includes the Maximum Authorized Speed (MAS) and temporary speed restrictions. The train control component enforces the requirements of the signaling.

There are a number of differing PTC systems being developed by suppliers with no universal system being available.

There are five primary PTC systems currently being developed or in operation are:

ATC – Automatic Train Control (cab signals and speed control)
ACSES – Advanced Civil Speed Enforcement System (usually applied as a vital overlay to ATC)
ETMS – Electronic Train Management System
I-ETMS – Interoperable Electronic Train Management System
ITCS – Incremental Train Control System

Each of the systems has succeeded in delivering the core requirements of PTC. There are also various Communications Based Train Control (CBTC) systems being installed or considered by many high-density transit systems that could satisfy the PTC requirements.

3. ACSES

During the early 1990's Amtrak began to work with the Federal Railroad Administration (FRA) to implement a maximum speed of 150 Miles per Hour (MPH) over the North-East Corridor (NEC). The Order of Particular Applicability for the NCE to supplement CFR 49 sections 235 & 236 with ACSES was issued on July 22 1998 for the entire NEC. The first revenue service under ACSES/ATC was from Stonington, CT to Canton Jct., MA in December 2000. This was expanded to New Haven, CT and Boston, MA in 2001 as well as to certain high speed tracks in New Jersey and Maryland.

ACSES is designed as an overlay to the pre-existing cab signal based ATC system in place on the NEC. The cab signal system provides protection of trains by forcing a following train down to 15 MPH prior to reaching an occupied block. The cab signal system does not enforce positive stops and does not enforce permanent or temporary speed restrictions.

ACSES works with the cab signal system and fills in by enforcing temporary and permanent speed restrictions and by enforcing positive stops at home signals displaying an absolute stop aspect. ACSES works through a combination of passive transponders, wayside processors, a radio network and an on-board processor that is connected to the onboard ATC system.

Transponders are placed along the track at select locations. Transponder locations are typically selected to surround interlockings and stations and to ensure that a train never travels more than 1.5 to 2.5 miles without passing a transponder. The transponders serve to correct for positional error in the location computations of the on-board processor and also carry fixed information pertaining to grades, civil speed restrictions and Positive Train Stop (PTS) locations among other information.

Information on temporary speed restrictions (TSRs) and the status of home signals is provided via radio messages. The onboard system receives information from the transponders as to how to address the radio network. The

on-board system then requests, through the radio network, the status of any home signals being approached as well as any TSRs in effect.

The onboard ACSES system tracks train speed, direction and location by counting wheel revolutions and from data read from the transponders. Direction is determined by the order that the transponders are read, location is known in absolute terms when a transponder is passed and is then computed by counting wheel revolutions from that point. Wheel wear can be checked and compensated for by comparing the computed distance between transponders to the actual distance when the transponders are passed. The system takes a large enough sample of these comparisons to ensure that there is a trend indicating wheel wear. Slip/slide is accounted for by adding a variable buffer of 5% of the distance from the last transponder to the braking computations.

3.1. ACSES Braking Calculations

ACSES works on a "distance-to-target" principle. The allowable speed for a train and a given distance from the braking target is given by either an equation or a lookup table. ACSES also computes a warning distance, a location where the train operator will receive a warning that the train is about to contact the braking profile. This warning distance is computed to provide a specified amount of operator reaction time at the train's current speed.

3.1.1. Train types

ACSES supports five different train types each with its own set of braking criteria. The train types are currently defined as follows:

Type A (or type 1) - High Speed Trainset with tilting

Type B (or type 2) - High Speed Electric Locomotive or High Speed Trainset without tilting or AEM-7 with Metroliner cars

Type C (or type 3) - HSEL, AEM-7, other passenger with Heritage fleet, commuter rail

Type D (or type 4) - Locomotive with Mail/Express

Type E (or type 5) - All freight operations

Each trainset operating under ACSES must be categorized as one of these types. The selected train type determines the brake rate that will be enforced under ACSES

If a train reaches the ACSES warning curve the engineer will receive an audible alert and the new target speed will be displayed. The warning curve is typically defined to provide the engineer with 8 seconds of reaction time before reaching the penalty curve but this number can be configured within the ACSES on board processor.

If a train reaches the ACSES penalty curve the ACSES system will initiate a brake application. When the train is once again below the ACSES warning curve the brake application can be released and the engineer can again take control of the train.

3.1.2. Grade handling

ACSES supports a single “worst case” grade for each braking target. This grade is programmed into the transponders and passed to the on board system via the transponder message identifying the braking location. This grade programmed into the transponder is referred to as the “controlling grade”.

Typically the controlling grade will be the worst case (largest downgrade or least steep upgrade) encountered from the starting point of braking to the point of protection (speed restriction or signal). Since there is only a single controlling grade for each braking point for all train types the start of braking must be determined using the worst case train type.

ACSES then uses an equated distance calculation to determine the required braking distance based on the controlling grade. This controlling grade can be set by ACSES engineers to account for unusual circumstances (very short grades, etc.)

3.1.3. Signal stop release

Because ACSES works in conjunction with the existing ATC system the ACSES system typically knows that an interlocked signal in front of the train has cleared when an upgraded cab code rate is received. In certain scenarios a code rate upgrade will not be received when an interlocked signal clears. In these cases the enforced stop for that signal is cleared through a radio message. When this happens there can be a delay in receiving this release which is dependent on the processing cycles of the wayside equipment.

3.1.4. Error handling

As previously stated the primary source of uncorrected error is from slip/slide of the wheels. Since track conditions are not known by the onboard processor a conservative estimate is made of the worst case error. The error is corrected every time the train passes over a transponder set. From this point until another transponder set is reached, the error used by ACSES at any given location is five percent of the distance traveled from the previous transponder set. The assumption is that the train could be 5% short of the target location or 5% beyond the target location based on whether the wheels slipped during acceleration or slid during braking.

There is also a fixed buffer enforced for signal stops. Amtrak is currently using 100 feet for this value but this can be configured. Other operators on the NEC have more stringent requirements for stopping a train, particularly when an interlocked signal is on the end of a platform or

there is a short passing siding. Requirements have been written for ACSES to reduce this buffer but there are practical limits in the system.

3.1.5. On-board display units

ACSES display units originally displayed to the engineer the current commanded speed. When the train crosses an ACSES warning curve an audible alert is sounded and the new speed is displayed. Some freight trains were equipped with an auxiliary display unit which displays the penalty speed above the current train speed to provide additional feedback.

The display units for ITCS have a Time to Penalty (TTP) indication that has been found to be very useful. The TTP indication displays the number of seconds before the train reaches the penalty curve and gives the engineers instant feedback if they are braking too steeply or not steeply enough. Amtrak has decided to add a TTP indication to their new display units. Other operators on the NEC are considering this. There is a concern by some that the TTP indication would be a distraction to engineers.

With the TTP indication an engineer could theoretically follow the ACSES penalty curve to the target at a safe interval of a couple of seconds without the uncertainty and potential over-braking that would otherwise occur.

4. RAILSIM ACSES IMPLEMENTATION

The RAILSIM simulation software suite has been used for many years to simulate high density train operations including operations on the NEC. RAILSIM is the most accurate rail simulation tool available and has been calibrated to a high degree of accuracy on many complex high density systems. Train performance and signaling in RAILSIM is designed to be flexible and easily customizable.

ACSES has been modeled in RAILSIM to allow simulation of the varying conditions expected when operating under ACSES. The following modifications were made to RAILSIM to support this.

4.1. Definition of ACSES Territory

A flag on the RAILSIM track segments allows the user to specify the boundaries of ACSES territory. Yards and other lead tracks can be marked as non-ACSES territory and train modeling in these sections will be performed as normal without ACSES.

4.2. Transponder Sets

RAILSIM allows transponder sets to be located anywhere along the track layout. Transponder messages can be encoded for each braking location beyond the transponder with a controlling grade and, for signals, a fixed setback distance to stop the train short. The signal setback is

initially set to a user-defined default value but can be changed for each individual signal braking location. The controlling grade is initially set to the worst case grade between the transponder set and the target braking location but can be re-set by the user for individual target locations.

4.3. Train Performance Options

Train performance in ACSES territory is determined by a number of user configurable parameters. The ACSES train type is selected first, this determines the ACSES enforced braking curves that the train must follow to avoid a penalty brake application.

The RAILSIM user would then select from a number of configurable parameters that determine how close an engineer would get to the ACSES penalty curve before beginning to apply the brakes, how aggressively the engineer would apply the brakes and, ultimately, how far from the ACSES target location the train will reach the enforced speed. This information would typically need to be determined by monitoring trains running under ACSES and observing how engineers react to the enforced braking curves.

The RAILSIM user can select whether a given train class is equipped with a TTP indication. If this is selected RAILSIM will allow a user to select a time to maintain from the penalty curve and the train will follow the curve at this distance.

4.4. ACSES Braking Curves

RAILSIM computes the ACSES warning and penalty curves using the same calculations as are performed by the ACSES on-board processor. Fixed and variable buffers are added to stopping locations based on transponder layout and user parameters. Braking locations and rates for each train are then computed based on a combination of these curves and the user defined train performance parameters.

4.5. Signal Clearing

When a signal clears, RAILSIM will not allow a train to proceed until either a cab code upgrade is received from the signal behind the train or a radio release time has expired. The radio release time can be either a fixed amount of time or can be randomized based on user entered parameters.

5. ACSES SIMULATION

5.1. Model

A sample model will be used for the simulations. The model consists of approximately 15 miles of 4 track cab signal territory. This was originally part of a larger real-

world model but grades and speed restrictions were added to this model, in close proximity in some cases, to test and demonstrate the effect of grades on the ACSES braking calculations. A mix of push-pull and EMUs is used for the various consists in the simulation.

A conceptual transponder layout was added to the model following the typical rules for locating ACSES transponders. Figure 1 shows the model used. The civil speeds are on the top with grades below. The triangles on the schematic indicate transponder positions.




Figure 1 - Sample ACSES model

5.2. Transponder Layout

A basic transponder lay out was put into the model such that a transponder would be encountered at the exit from every interlocking as required by ACSES so that the system knows what track the train is on. Transponders were then placed to ensure a train would travel no more than 1.5 to 2.5 miles before reaching a new transponder set. This is a recommended practice for ACSES design. In some cases transponder layout was adjusted so that transponders would be placed just before platforms to increase the accuracy of ACSES braking curves for signals in close proximity to a station.

No other design considerations were taken into account in the transponder layout. In a normal ACSES design other considerations would be taken into account such as changing grades near a braking location.

5.3. Train Performance

ACSES is a relatively new system and there is not a lot of empirical evidence for how the engineers would operate a train under ACSES. One would expect, and anecdotal evidence shows, that engineers operating under the system for the first time would tend to start braking after receiving the alert and brake at a rate close to what they would use under normal operations without ACSES. These rates will tend to be more aggressive than what ACSES is enforcing, particularly in the case of changing grades, and the trains

would end up braking well short of the enforced target location.

The same anecdotal evidence also shows that with training and time engineers would learn to run their trains closer to what the ACSES system is enforcing. Engineers would learn where the ACSES warning curve is and would likely begin braking prior to receiving the alert. They would then learn to brake at a rate where they would reach the enforced speed very close to the ACSES enforced target location.

Not having hard evidence of operator performance in the presence of ACSES we will model a “worst case” and “best case” scenario. The worst case scenario will account for engineers learning to brake prior to reaching the warning location but still braking at too aggressive a rate. These scenarios will be compared to a scenario without ACSES installed or operating.

5.3.1. Scenario 1 – no ACSES

This scenario will model operations without ACSES installed. Engineers who have been operating trains over a particular railway for many years will know where they need to start braking and how aggressively they need to brake to satisfy speed restrictions and stop for signals. Engineers will likely reach their target speeds very close to curves since there is no enforcement mechanism. At signals they will leave some small buffer in where they stop but likely no more than a few feet.

This scenario will use braking rates for signals, speed restrictions and station stops that have been collected and calibrated over many years of observing and simulating real train operations in the NEC.

5.3.2. Scenario 2 – worse case

Engineers will quickly learn where they need to start braking and will likely brake a little before reaching the ACSES warning curve. This scenario will model this behavior without the engineer adjusting properly for the ACSES enforced brake rates. This is considered a worst case scenario and should not be encountered if engineers are sufficiently trained before ACSES is put in place. There may, however, be some very conservative engineers who will initially run this way. Trains will end up braking well short of the ACSES target locations.

5.3.3. Scenario 3 – best case

This scenario will model the best case where engineers learn where the ACSES warning curve and target locations are and initiate a smooth braking curve starting just prior to the warning location and ending very close to the ACSES target braking location.

5.4. Simulation Results

5.4.1. Scenario 1 – no ACSES

This scenario started with a current real-world operating plan over the selected segment. Since grades and speed restrictions were added to the model the schedule had to be adjusted for the new train running times. The simulation was optimized based on the following requirements

1. No train should be later than 6 minutes
2. Non-revenue trains only need to reach their end destination on or ahead of schedule. Delay is not a factor
3. Delay for revenue trains was minimized to the greatest extent possible

Table 1 shows the system lateness for this simulation. Table 2 shows the lateness distribution.

Average Lateness: All Trains	Average Lateness: Late Trains	Latest Train
0:49	1:24	0:05:04

Table 1 - System Lateness

Lateness Distribution	Train Count	Percent Of Trains
0:00 to 1:00 minutes early	28	41
0:01 to 1:00 minutes late	21	30
1:01 to 2:00 minutes late	9	13
2:01 to 3:00 minutes late	6	8
3:01 to 4:00 minutes late	2	2
4:01 to 5:00 minutes late	1	1
5:01 to 6:00 minutes late	1	1

Table 2 - Lateness Distribution

Trains merge at the west end of the model from 4 tracks down to 2. This means that local trains running on the outer tracks need to merge with express trains on the middle tracks. There are also a number of non-revenue trains that traverse the network from the west end to the east end and turn around either at the crossovers in the middle of the model or at the east end of the model and go into service at that point.

Priorities at these merge and conflict points were generally given to the train with the earliest schedule time beyond this point, to allow trains to make their schedule. The operating plan has specific slots reserved for these turning and merging operations and trains generally fit well into these slots without causing an inordinate amount of delay.

5.4.2. Scenario 2 – Worse case

In this scenario the trains will brake prior to reaching the ACSES warning curve and brake more aggressively than is necessary to satisfy the ACSES enforcement. The intent is just to find any potential issues or schedule adjustments that might need to be made so no attempt was made to re-prioritize or re-schedule trains for this scenario.

Table 3 shows the new system lateness and Table 4 shows the lateness distribution for this scenario.

Average Lateness: All Trains	Average Lateness: Late Trains	Latest Train
1:46	1:56	0:08:20

Table 3 - System Lateness

Lateness Distribution	Train Count	Percent Of Trains
0:00 to 1:00 minutes early	6	9
0:01 to 1:00 minutes late	19	28
1:01 to 2:00 minutes late	22	32
2:01 to 3:00 minutes late	8	12
3:01 to 4:00 minutes late	7	10
4:01 to 5:00 minutes late	3	4
5:01 to 6:00 minutes late	1	1
6:01 or more minutes late	2	3

Table 4 - Lateness Distribution

5.4.3. Further analysis

Figure 2 shows an express train running under ACSES and compares the ACSES run to the non-ACSES run. Figure 3 shows the same graphic for a local train. As can be seen, ACSES has a much greater effect on the express trains than the local trains since the stopping pattern for the local

trains masks much of the effect. Table 5 shows the difference in running time with and without ACSES for both the local and express trains.

The two trains later than 6 minutes were due to conflicts that arose when a local train set its route at the merge point before an express train that should have been in front of it. Giving the express train a priority at this merge point resolved this issue but the local train then received an additional delay. To resolve both issues, the schedule needs to be modified to account for the new running times.

Service	Running Time		Difference
	No ACSES	ACSES	
Local	24:58	25:35	00:37
Express	17:11	18:16	01:05

Table 5 - Comparison of running times




Figure 2 - Express Train




Figure 3 - Local Train

5.4.4. Scenario 3 – Best case

Scenario 3 simulated the best case where the engineers all operated very close to the ACSES enforced braking curves and reached target speeds very close to the ACSES target locations. This scenario saw far less disturbance in

operations. As can be seen in Table 8, the local trains were very close to the non-ACSES running times and the express trains were delayed less than 30 seconds. The average lateness actually came down from the non-ACSES run. This is likely just coincidental as conflicts may have been reduced when the running times of the trains changed slightly. The same effect could be had by just shifting the train schedules in the non-ACSES simulation.

There are fewer earlier trains but no train is more than 5 minutes late. No additional conflicts were seen in this run.

Average Lateness: All Trains	Average Lateness: Late Trains	Latest Train
0:53	0:58	0:04:39

Table 6 - System Lateness

Lateness Distribution	Train Count	Percent Of Trains
0:00 to 1:00 minutes early	9	13
0:01 to 1:00 minutes late	42	62
1:01 to 2:00 minutes late	10	15
2:01 to 3:00 minutes late	2	3
3:01 to 4:00 minutes late	4	6
4:01 to 5:00 minutes late	1	1

Table 7 - Lateness Distribution

Service	Running Time		Difference
	No ACSES	ACSES	
Local	24:58	25:05	00:07
Express	17:11	17:38	00:27

Table 8 - Comparison of running times

5.4.5. Other considerations

5.4.5.1. Signals at the end of platforms and sidings

This model contained some interlocked signals at, or very close to, the end of the platforms. ACSES as configured in this simulation would not allow the trains to berth at the platform while these signals were at stop. The RAILSIM

network simulator would normally defer requesting these routes until the trains were almost ready to leave the platform but in this case we had to allow trains to set their routes through the interlocking before entering the station.

This is a situation likely to be encountered in the real world and is currently under study. A similar problem can be encountered with train-length sidings where the train needs to approach very close to a signal in order to clear the fouling point of a switch behind.

In this case setting these routes before entering the station only caused minor delay in a couple of cases. In real world situations this will need to be addressed by a combination of:

1. Adding transponders to reduce the error due to slip/slide
2. Reducing the fixed buffer for signal stops
3. In certain situations the ACSES enforced target location can be set beyond the signal but prior to any switch or fouling point.

RAILSIM supports simulation of all of these scenarios.


5.4.5.2. Additional transponders

Some of the running time delay encountered in these simulations can be reduced by adding additional transponder sets. Figure 4 illustrates one case where the transponder layout is not optimal. The civil profile is shown in green, the ACSES penalty curves are shown in red and the warning curves in blue. The transponder locations are shown in magenta.

The transponder just before the speed restriction is too close to have any impact since the train will be forced down to speed before reaching it. The previous transponder is controlling the speed but since it is so far back the slip/slide error forces the train down to speed well before reaching the speed restriction.

If an additional transponder set was added close to the start of braking for this speed much of this error would be eliminated.

**Figure 4 - Computed ACSES
braking curves**



While this will necessarily have some negative impact on operations, with careful planning this impact can be minimized.

This is just an example of the design decisions that can have a substantial impact on the operation of ACSES. Reporting in RAILSIM is currently being enhanced to provide more insight into the accuracy of ACSES braking curves under a specified design.

6. CONCLUSIONS

The installation and transition to any PTC system needs to be well thought out and carefully designed. Training of engineers is an important component of a successful transition to PTC based operations and certain design decisions can have a significant impact on future operations.

The most likely outcome of a transition to ACSES control would be somewhere in between the best case and worst case scenarios described here. Initially performance may tend closer to the worst case as engineers learn a new way of running their trains. There will be a mix of conservative and more aggressive engineers who will react differently to a new way of braking under ACSES. With training and experience engineers should be able to run trains under ACSES without a substantial increase in run times.

There are also many additional design considerations that need to be studied for each individual ACSES installation. Transponder layout can have a significant impact on train performance under ACSES and needs to be studied in the context of grades and braking profiles at any given location. Error tolerance is also a very important consideration and will vary according to the proximity of interlockings to stations, the length of sidings and the distance between interlockings. Requirements have been written for planned ACSES implementations to tighten these error tolerances but the success of this effort remains to be seen.

ACSES has unique characteristics, as do all PTC systems, but the nature of the potential issues in ACSES is not unique. All PTC systems have sources of error that need to be accounted for and all will take a conservative safety focused approach to enforcing speeds and positive stops.