SAND--98-8524C

COORDINATED TRAIN CONTROL AND ENERGY MANAGEMENT CONTROL STRATEGIES

CONF-980002-

RECEIVED

MAY 11 1998

OSTI

S.P. GORDON
Sandia National Laboratories
MS 9201, P.O. Box 969
Livermore, CA 94551-0969

D.G. LEHRER
Bay Area Rapid Transit District
212 - 9th Street, P.O. Box 12688
Oakland, CA 94604-2688

ABSTRACT

The Bay Area Rapid Transit (BART) system, in collaboration with Hughes Aircraft Company and Harmon Industries, is in the process of developing an Advanced Automatic Train Control (AATC) system to replace the current fixed-block automatic system. In the long run, the AATC system is expected to not only allow for safe short headway operation, but also to facilitate coordinated train control and energy management. This new system will employ spread spectrum radios, installed on-board trains, at wayside locations, and at control stations, to determine train locations and reliably transfer control information. Sandia National Laboratories has worked cooperatively with BART to develop a simulator of the train control and power consumption of the AATC system. We are now in the process of developing enhanced train control algorithms to supplement the safety critical controller in order to smooth out train trajectories through coordinated control of multiple trains, and to reduce energy consumption and power infrastructure requirements. The control algorithms so far considered include (1) reducing peak power consumption to avoid voltage sags, especially during an outage or while clearing a backup, (2) rapid and smooth recovery from a backup, (3) avoiding oscillations due to train interference, (4) limiting needle peaks in power demand at substations to some specified level, (5) coasting, and (6) coordinating train movement, e.g. starts/stops and hills.

INTRODUCTION

In order to support future train operations, the Bay Area Rapid Transit system is faced with a need to upgrade the existing train control system to enable more trains to operate on existing track infrastructure. The current fixed-block automatic control system is capable of operating trains at a crush headway of approximately 150 seconds. As demands on the system increase, shorter headways will be required in order to postpone the need to build expensive additional track infrastructure. The Advanced Automatic Train Control system under development at BART promises to achieve shorter headways by locating trains to within 15 feet and defining moving control blocks based upon the train locations.[1] Following distances will therefore be constrained only by safety concerns, while the large position uncertainty of a fixed block system will be removed.

The AATC will employ the Enhanced Position Location Reporting System (EPLRS), a spread-spectrum radio ranging technology developed by Hughes and capable of simultaneous train tracking and communication.[2] Trains will communicate from on-board radios through a network of wayside radios to station computers, which will control the trains in local areas through speed and acceleration commands updated every 0.5 seconds. Both train locations and speed commands will be more finely resolved than in the current system, allowing for more precise control.

In order to test and refine the AATC system before implementation, the Train Control Simulator (TCS) was developed at BART to accurately simulate the motion of trains.[3] The TCS then became the central part of a larger AATC Train and Power Simulator (ATAPS), which supplemented the safety-critical control system in the TCS with a non-vital controller, allowing for the addition of enhanced control algorithms for coordination of multiple trains.[4] ATAPS also contains a system

MASTER

DISCLAIMER

DISCLAIMER

traction power model, which allows analysis of power-related issues, such as power consumption and low voltages. The first half of this paper will describe the simulator, and demonstrate its capabilities.

The simulator is now serving as a testbed for enhanced control development. The control logic of AATC is resident on wayside computers, which calculate speed and acceleration commands for all trains within their control zones. Thus, it is possible to coordinate the motion of multiple trains in order to improve service reliability and passenger comfort, while reducing energy infrastructure and usage costs. The second half of the paper will outline the types of enhanced control that we have considered and the associated benefits which may be achieved.

THE SIMULATOR

Development of the Train Control Simulator (TCS) began at BART in the early nineties to explore interest in a moving block train control system. The first versions of the simulator were rudimentary – designed specifically to evaluate the advantages of a moving block control system compared to a fixed block system in terms of headway. Eventually the data provided by these simulations led BART to go forward with the AATC project.

The new train control approach brought with it the need for more sophisticated simulations, so the TCS began a phase of ongoing modification and upgrade. Newer versions provided the capability to simulate different braking and acceleration control algorithms, to estimate run times, and to test changes made to the vehicle control code. The latter has developed into a key technique in the validation of such modifications, allowing for more exhaustive testing than might otherwise be possible.

More recently, the TCS has been integrated with a traction power model developed at Sandia National Laboratories to create the AATC Train and Power Simulator (ATAPS). This simulator is capable of modeling the motions of trains traveling in both directions on a single line of the BART system, as well as predicting the state of the power system at any given moment. In its current form, the trains are modeled over the entire length of a line, whereas the power model extends only over a subsection of that line (e.g. the trains run from Daly City to Fremont and back, while the power system is modeled from Oakland West to the 24th Street San Francisco substation).

Train Control Simulator

The BART Train Control Simulator is unique in that it incorporates the actual vehicle-borne control code into the train motion calculation. This produces extremely accurate, high resolution simulation data, down to 36 millisecond intervals. With the aid of this precise simulation, BART has further optimized its existing vehicle control system, which had gone unmodified for nearly a decade.

Although the use of vehicle-borne code makes the simulator BART-specific, the design is modular enough to allow alternate or more generic train motion code to be swapped in. Future versions of the simulator may be able to select from multiple vehicle types as easily as clicking on a pull-down menu.

Features

The TCS has many user definable parameters. It is capable of simulating conventional track circuit-based control systems as well as moving block systems, or a mixture of both. Multiple trains may be simulated simultaneously, each with its own configuration; length, weight, performance profile, and speed selection algorithm are all user definable.

The type of output data and its resolution in time are also user definable. Reports may be generated which log events, such as track circuit occupancies, switch and gate movement, as well as relative position graphs depicting the proximity of trains to one another. All output data is available in ASCII text format.

Design

The majority of the simulator is coded in C and C++. The vehicle control code modules are coded in 80x86 assembler. The program is constructed in modules which functionally represent the train control system elements that they simulate. For example, there are C++ classes which represent the track, the interlocking control system, the station computer and the trains themselves. This modularity allows different subsystem designs to be tested by swapping classes within the simulator.

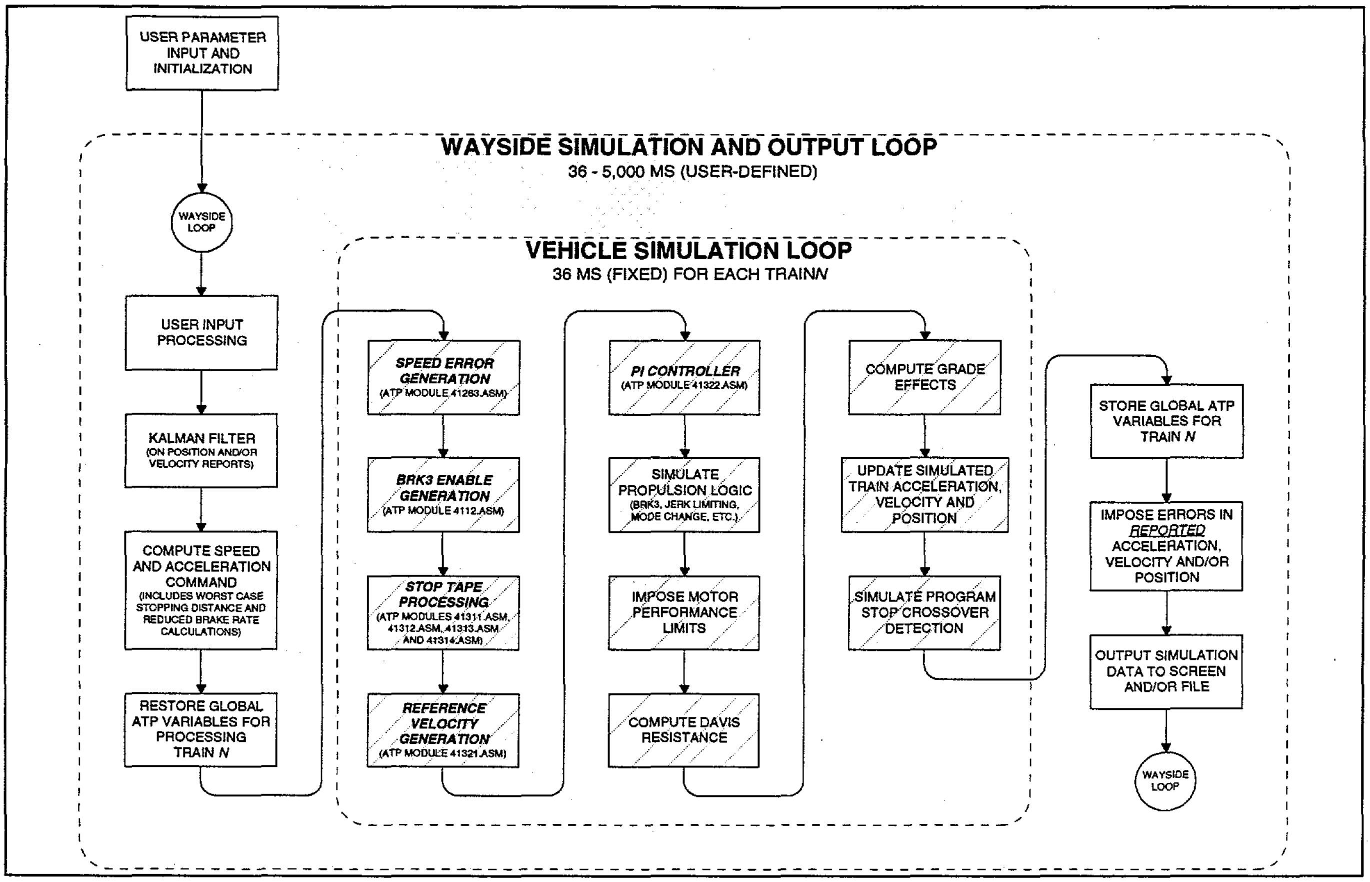

Figure 1. Train Control Simulator flow chart

Figure 1 depicts program flow among the major subsystems within the simulator. The external wayside simulation loop includes all user input functions, speed code generation and output logging. The timing of this loop is adjustable, depending on the control interval and data resolution required.

The internal vehicle control loop includes the actual vehicle-borne control code and motion equations. This loop operates at fixed 36 millisecond intervals, to replicate the design of the vehicle code. Future enhancements to the simulator may include the ability to adjust this time interval as well, allowing more complex track layouts and more trains to be simulated in less time.

Accuracy

The accuracy of the simulator has always been a paramount design concern. To date the TCS has been validated against system wide revenue service data collected by BART's Central Control System computer, and against data collected during the first phase of AATC system testing. Both have indicated that the simulator precisely models the control and movement of BART trains.

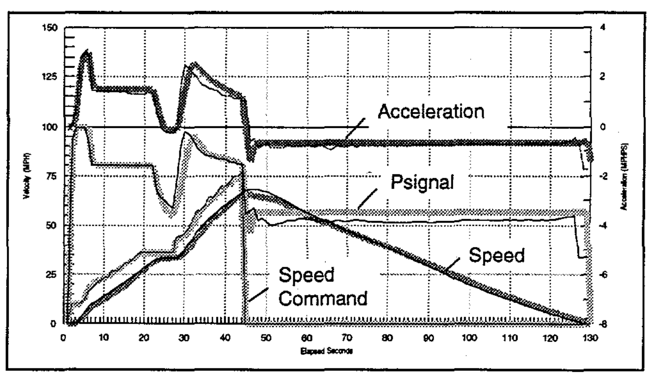

Figure 2. TCS validation data

In Figure 2, the thin lines represent the speed and acceleration of an actual train under AATC control during Phase 1 testing on BART's Hayward test track. The thick lines overlapping the thin lines show a simulated train traversing the same section of track. Also shown are the command speed and the on-board propulsion signal (Psignal). The data resolution is approximately 500 milliseconds. As this figure shows, the train simulator accurately models the actual train motion.

Discoveries

Many potential problems have been revealed using the simulator, some far sooner than they might otherwise have been discovered. Early on in the AATC project it was found that trains were routinely undershooting their target brake rates; that is, trains were braking too hard, often by as much as 10%. Investigation revealed that a software module called the Proportional-Integral (PI) Controller, which is part of the onboard vehicle code, contained parameters which were not set optimally. In addition, the code contained a sequencing flaw.

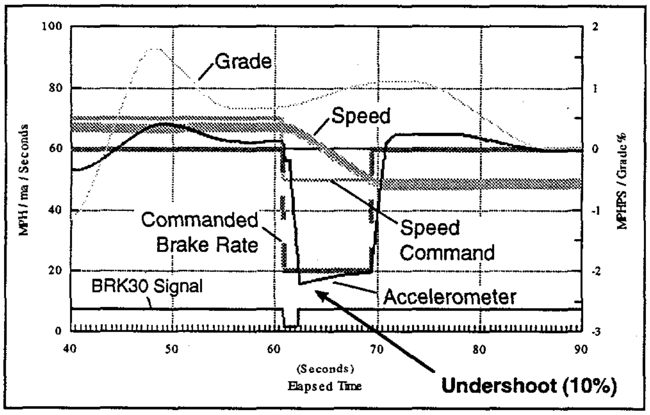

Figure 3. Simulation of original PI Controller

Figure 3 shows a simulated train making the transition from propulsion into braking with the old control code. It is clear that the train's acceleration rate, shown in the accelerometer plot, initially undershoots the commanded rate. Figure 4 shows the same scenario with the new control code. The simulator facilitated the discovery and correction of this problem. Over time, such small optimizations can add up to substantial savings.

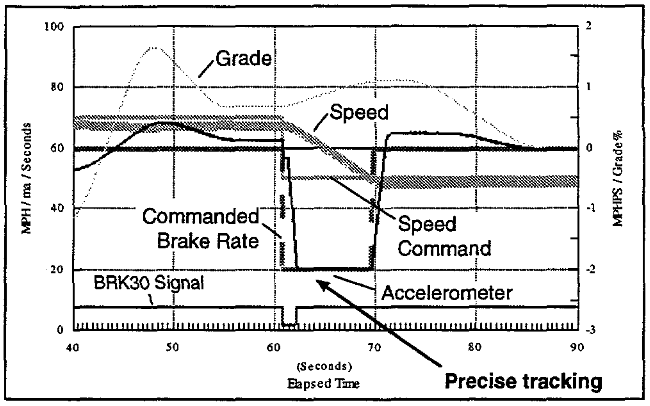

Figure 4. Simulation of revised PI Controller

Later in the AATC project, the simulator revealed that a cyclical behavior could result if two trains stopped close together attempted to accelerate up to speed. During acceleration, the following train would be forced to repeatedly stop accelerating, and sometimes begin braking, in order to maintain a safe distance from the train ahead. This phenomena is rooted in the calculation of safe stopping distance, which is proportional to the velocity squared. If the following train accelerates at the same rate as the lead train, its expanding stopping distance quickly overtakes the lead train, and it must stop accelerating. Once its following distance increases sufficiently, acceleration resumes, and the process is repeated. An example of this is depicted in Figure 5. Based on the forewarning provided by the simulator, enhanced acceleration control algorithms were developed to prevent this oscillatory scenario.

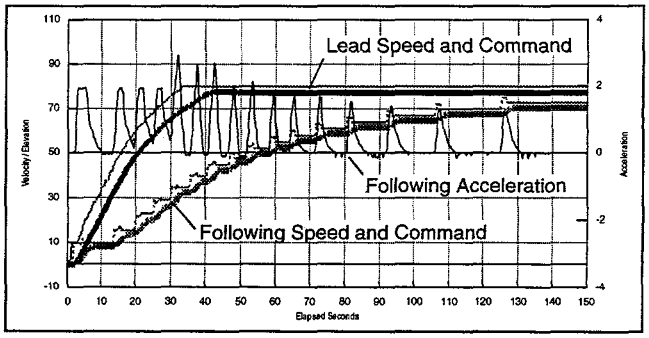

Figure 5. Impeded Following Train Acceleration

The simulator has also proven to be an invaluable tool for the prediction of overall system performance capabilities. Using the simulator to model the new AATC moving block control system and comparing the run times with simulated track circuit-based control, an estimate of potential run-time improvement was derived. This improvement indicated that entire trains may be eliminated from future schedules, thus reducing operational and capital costs. In addition, the simulator indicates that AATC will be capable of much shorter train headway, allowing for faster recovery after delays.

AATC Train and Power Simulator

The Train Control Simulator described above comprises the core of the AATC Train and Power Simulator (ATAPS), which also contains a traction power model for simulation of system power flow, as well as a non-vital controller for adding enhancements to the safety-critical vital control. As the simulator runs, it generates an output file which records train trajectory and power consumption data and substation power production data as functions of time. These data, along with the output files generated by the TCS itself, allow for a detailed analysis of the impacts of various control algorithms. In addition, the Benefits Assessment Module (BAM) evaluates a small set of overall power- and control-related metrics, such as the number of voltage sag events, so that the user can easily assess the overall impact of an algorithm without looking at the detailed output data. We believe the ATAPS simulator provides a powerful testbed for development and testing of novel train control algorithms in the context of both train motion and power consumption.

The models of traction power and of the non-vital controller, as well as the Benefits Assessment Module, are all interfaced to the TCS with a single function call. By using this function as a container for other functions that do the "real work," the additional ATAPS code modules may be maintained and updated with little impact on the TCS itself. This structure allows BART and Sandia, in our collaborative relationship, to easily update our separate sections of the code independently. In addition, a power model capable of modeling more complex track geometry than a single line could easily be substituted for Modrails at some future date. The only exception to this single-interface philosophy is the train motor model (PMAM, described below), which is incorporated directly into the TCS.

Traction Power Simulator

The Sandia traction power model, Modrails (Model of DC Rail Systems), uses the location and power consumption or regeneration of each train at a given moment to calculate the voltage at each train and the power being produced by each substation. The system solution is found in the steady state, so Modrails does not support studies of AC instabilities in the power network. The utility of the model is in evaluating the severity of voltage sags and the usage of regenerated traction power.

The power demand of each train is first calculated by the train Power and Maximum Acceleration Model (PMAM) within the TCS. It calculates the power consumed or regenerated by a train, and the maximum possible acceleration which the motors can provide given the current train state. This function not only impacts power calculations, but also limits train trajectories to physically realizable accelerations.

Given the calculated power consumption and location of each train on the line, Modrails translates all infrastructure and train information for a linear section of track into a DC electrical circuit. Incorporated into this model are substation, crossbond, and gap breaker locations derived from the BART track plans. Crossbonds are connections between the running rails in the two directions of travel, and gap breakers are connections between the two powered contact rails. Each train is treated as two separate power sinks (or sources) consuming (or regenerating) half of the total train power, one located at the head of the train, and the other at the tail. This avoids overestimating voltage sags by realistically distributing the power load. Train voltages are limited to a maximum of 1150V during regenerative braking. Power is allowed to flow out of substations onto the third rail, and from running rails into grounds, but not in the reverse directions. Rail resistance may vary in discrete sections, which allows accurate modeling of the presence of low resistance contact rail in some sections of track.

Non-vital Algorithm Simulator

In addition to simulating the baseline AATC system, ATAPS allows the user to implement non-vital train control algorithms in order to develop, test and refine enhanced control strategies. In the implemented system, the enhanced algorithms will reside in a non-vital computer which will communicate suggested train control commands to the vital station computer. The algorithms will use information about the state of the system received from the vital computer and through an ethernet connection to neighboring station controllers to calculate the suggested commands. Communication between the vital and non-vital computers will occur once every 0.5 second command cycle.

Sample Output Data

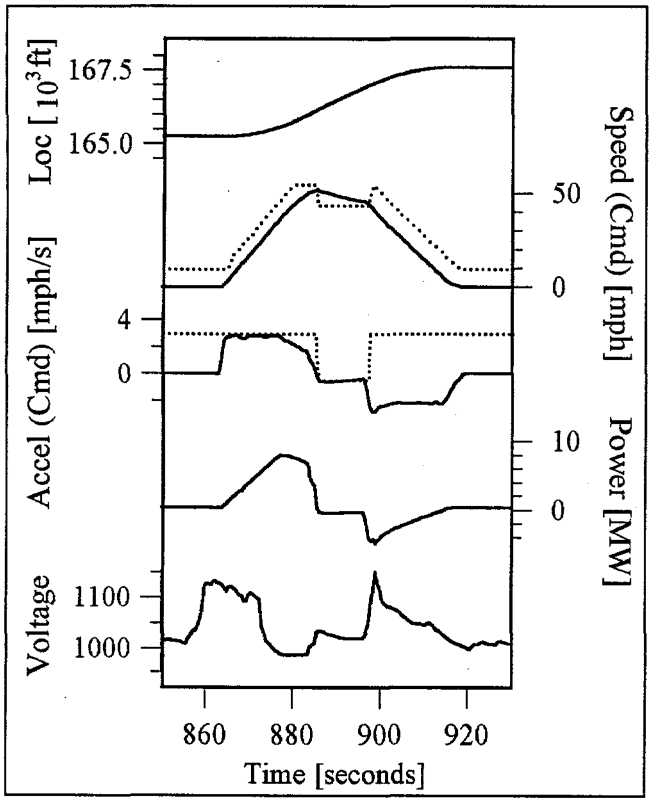

Figure 6. Train-related output data

Figure 6 contains the train-related output data produced by ATAPS during a typical run for one of the simulated trains. Train location, speed, command speed (dashed), acceleration, command acceleration (dashed), power consumption, and voltage are shown as functions of time. This train begins at a station, accelerates up to speed, and then decelerates for the next station stop. While the train is accelerating, power is consumed and the train's voltage drops. When the train is regeneratively braking, the voltage floats up to a maximum of 1150V. The commands shown do not match the trajectory during the station-stop, because the final braking for stations is controlled on-board rather than by AATC commands from the wayside station computer. Ultimately, it is expected that the entire train trajectory including station stops will be commanded from the station computers.

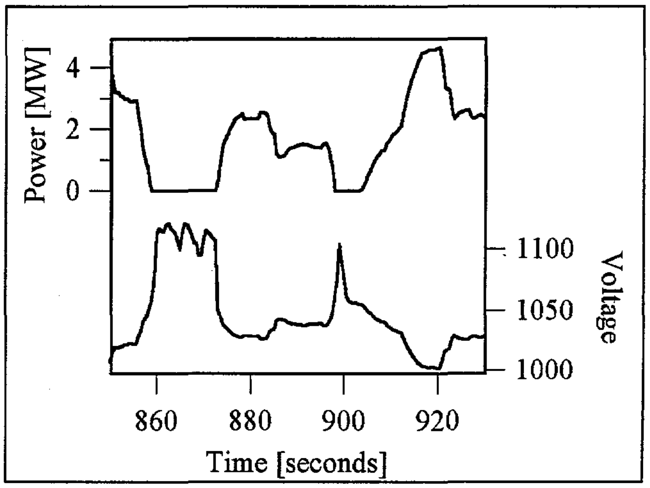

Figure 7. Substation-related output data

Figure 7 contains data from the same simulation run for one of the power substations. The power produced by the substation as a function of time, as well as the substation voltage, is shown. Power may only flow out of the substation, so power is always positive. When trains are regenerating nearby, the power drops to zero, and the voltage can float up well above the nominal 1050V.

ENHANCED CONTROL

The ultimate goal of enhanced train control is a system optimized with respect to a well developed and complete cost function. This function would represent the relative value of such things as trip time, energy and power usage, delay time, rider satisfaction, etc. Although such global optimization is a worthy goal, we believe that it is important to begin tackling the problem of enhanced train control by trying to solve more localized and well-defined problems. Therefore, we are attempting to design a few control algorithms which solve more specific problems, and assessing tradeoffs in their impacts on various important metrics. Once some intuition has been developed in this way, it may be possible in the future to design a more globally applicable algorithm.

From a long list of possible control objectives, we are concentrating on the five objectives shown in Table 1. For each power- or service-related objective, we have listed the relevant metric with which to measure it. Many of the algorithms discussed below primarily target service-related objectives such as a smoother ride for improved passenger comfort. However, a smoother ride typically corresponds to a reduction in unnecessary acceleration and braking cycles, which in turn reduces peak power consumption and overall energy usage. Thus, algorithms which target improved service often have energy-related benefits as well.

Objectives &	*Metrics*
Power-related	
Avoid energy infrastructure costs	Low train voltages
Reduce overall energy usage	Energy consumption
Service-related	
Improve service reliability	Train-minutes of delay
Improve passenger comfort	Smoothness of service
Minimize trip time	Trip time

Table 1. Algorithm objectives and metrics

The algorithms described below are in various stages of development. In some cases, they are only conceptual, and others have been implemented, and have begun to be tested and refined. This discussion represents an overview of the algorithms which we are pursuing, and the types of benefits that we expect from each.

Low Voltage Avoidance

The objective of this type of algorithm is to control multiple trains in an area to prevent low voltage events. When several trains which are close together demand power at the same time, if there is insufficient power available nearby due to either track geometry or an outage, then the voltage at the trains can drop precipitously. If the voltage drops below 750V, the motors presently on BART trains must shut down in order to avoid damage from excessive current flow. Even with motors that do not shut down, it is inefficient to allow severe voltage sags, as low voltages typically correspond to large power losses to heat in the rails.

An algorithm which prevents the voltage at each train from dropping below some minimum value by limiting power consumption can lead to large benefits in avoided energy infrastructure costs, in addition to reducing energy consumption. Power substations on the BART system are designed to maintain all train voltages above 750V even during times of heavy traffic with a substation outage. Thus, the power system is required to be extremely robust for normal operating conditions. If, on the other hand, a control algorithm were in place which could maintain train voltages above 750V by reducing speed and/or acceleration commands of lower priority trains when necessary, then it may be possible for the system to operate with a somewhat more modest power infrastructure. If this goal can be realized, then it may defer tens of millions of dollars of investments in additional substation power capacity.

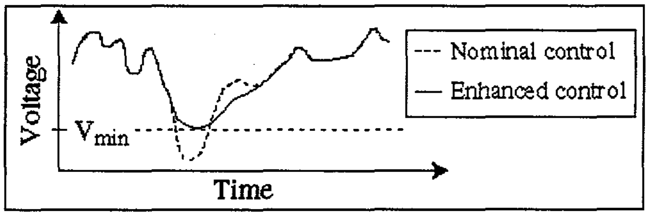

Figure 8. Maintain train voltage

Figure 8 shows schematically a train's voltage as a function of time with and without this type of enhanced control. In order to maintain the voltage of all trains above V_{min}, which would be somewhere in the neighborhood of 800V, it is necessary to predict train voltages based upon the trajectories of all nearby trains, and then allocate the available power in such a way as to maintain the voltage at all trains while minimizing the impact on the schedule. This algorithm, with such a high payoff, is not surprisingly very difficult to achieve. Train voltage is a non-linear function of power demand, which makes it difficult to predict quickly. In addition, power consumption can rise quickly enough to take the voltage from a comfortable range to well below the desired minimum almost instantly. Thus, it is not sufficient to measure or calculate train voltages and react as the voltage drops too low; rather, potential problems must be recognized before they materialize.

Rather than employing a slow but accurate system model such as Modrails to predict voltages, we are attempting to employ neural network technology to estimate voltages. The simulator is being used to train a neural network to predict low voltages. The network will then be integrated into an algorithm, which will reduce speed and acceleration commands when necessary to avoid sags. As a future enhancement, train voltages could be measured real-time, and this data used to continue to train and refine the neural net to improve its accuracy, and thus to increase the effectiveness of the algorithm over time.

A similar algorithm could be used to prevent needle peaks in power demand at substations. At the moment, these do not present a large cost, as substations are capable of providing high power for short intervals, and energy demand charges are presently for power averaged over half hour intervals. However, some energy infrastructure costs may be deferrable if power spikes could be avoided. In addition, if instantaneous demand charges become a reality in the future, then power demand spikes may become costly.

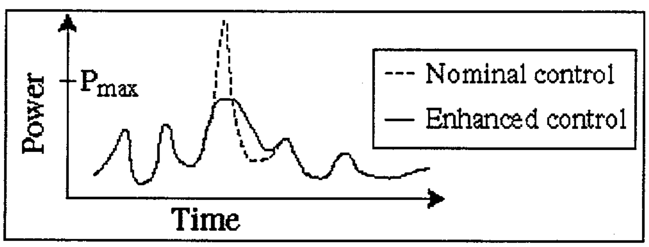

Figure 9. Limit peak substation power

A schematic of substation power as a function of time with and without an algorithm which caps power demand at P_{max} is shown in Figure 9. This algorithm will require a logic much like that for low voltages, but will be somewhat easier to implement because peak substation power is less strongly non-linear, and therefore easier to predict, than train voltage.

Interference Management

The objective of this type of algorithm is to avoid oscillatory brake/acceleration cycles due to interference, and, in general, to smooth the trajectories of closely-following trains. This is achieved by maintaining sufficient following distance to prevent unnecessary braking. The payoff from this algorithm comes from reduced energy costs and improved passenger comfort.

When a train follows very closely behind another, it tends to alternately accelerate and then brake as the two trains go over hills, or as they accelerate and brake between stations. As a train travels, it's predicted stopping distance can increase if the average grade in front of the train becomes more down-sloped. If that train is traveling as close as possible to the train in front of it, then it will have to brake if its stopping distance increases. This type of "interference" is predictable and preventable. Similarly, as two trains which are close together accelerate up to speed, the stopping distance of the rear train increases due to its increased momentum, and again oscillations can result. This type of interference was discussed in the TCS section above. A trajectory exhibiting interference during acceleration is shown in Figure 10 under "nominal control."

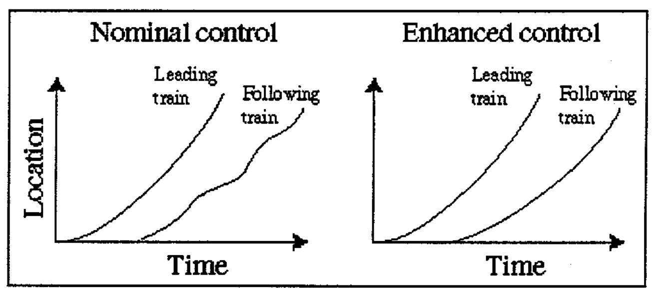

Figure 10. Smooth interfered operations

Although this behavior only becomes apparent during off-normal train operations when trains are abnormally close together, such as after delays, it is wasteful of energy, uncomfortable for passengers, and, most importantly, preventable. In addition, this "off-normal" condition will become more the rule than the exception as the scheduled headways become shorter to meet additional demand. Through enhanced control, the predicted stopping distance as a function of location and speed may be used to maintain the required following distance at all times, as shown schematically in the figure.

In addition to interference due to changes in stopping distance, avoidable braking may also result when trains are close together in a region with closely-spaced stations. If a train is stopped in a station, and another train approaches, the second train will begin to brake early to stop short of the station. If the train in the station then pulls away, the rear train may accelerate briefly before stopping for the station. This sequence may then be repeated at each station along the line if the stations are close together, as is the case in downtown San Francisco.

A relatively simple algorithm is capable of removing this type of interference. If a train is braking before a station because a train is stopped there, and it is then freed to accelerate by the stopped train pulling out of the station, then it should only accelerate if it is necessary in order to reach the station or if it will add excessively to trip time to remain in braking. Otherwise, it should continue braking at a low rate until the station stop.

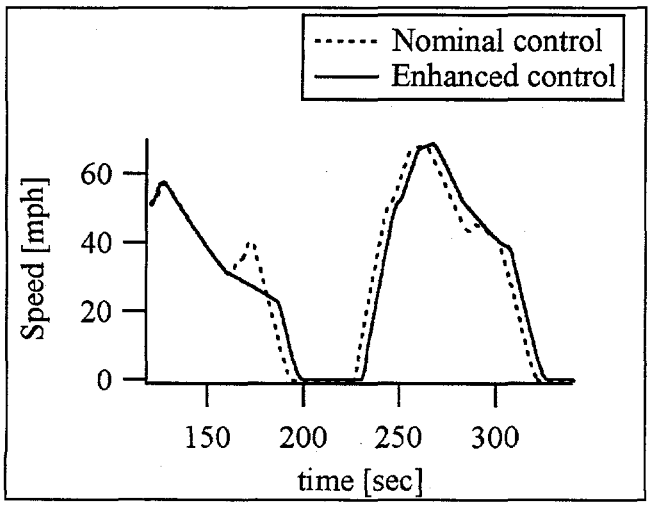

Figure 11. Smooth interfered station approach

Figure 11 shows an example of a train trajectory exhibiting interference before two consecutive stations. The velocity of the interfered train is shown as a function of time, as calculated by the simulator with and without enhanced control. At the expense of a few seconds of trip time, the enhanced trajectory is smoother, again saving energy and improving passenger comfort. The trip time increase shown here may be reduced by enforcing a more strict time restriction, which in this case would cause the train to accelerate to the first station, but to remain braking to the second station, where very little trip time is added.

Backup Clearing

Backups can lead to an extreme form of interference, where a line of trains sit one behind the other outside of a station, and the line moves forward one train length at a time as the trains pull into the station one after the other. This behavior leads to spikes in power demand, as the trains repeatedly accelerate and then brake to a stop. In addition to the resulting frustrating ride and the waste of energy, low voltages may result if sufficient power is not locally available for multiple accelerating trains. Although a low-voltage-avoiding algorithm may take care of this last problem, it is possible to use a more simple algorithm to smoothly and efficiently recover from a backup, while avoiding simultaneous accelerations which may cause low voltages.

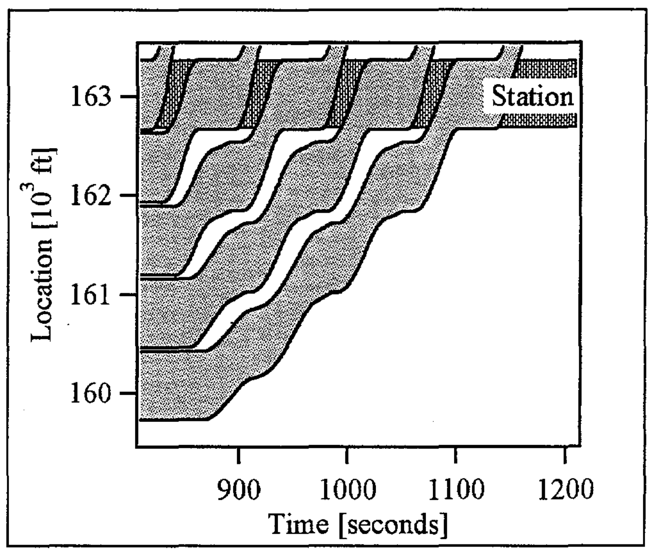

Figure 12. Nominal control after backup

Figures 12 and 13 show simulation runs of backups caused by a delay at a station. The full length of each train is shown as a shaded region along the location axis. The trajectory is flat when a train is stopped, and sloped when it is in motion. The line of four trains behind the stopped train moves through the station with nominal control in the first figure, and with a backup-clearing algorithm in place in the second figure. Under enhanced control, each train accelerates twice, once to come up to a calculated speed, and once to quickly pull into the station when it is clear. The trains are spaced out a bit more under enhanced control in order to avoid interference just before the station, but this is not required.

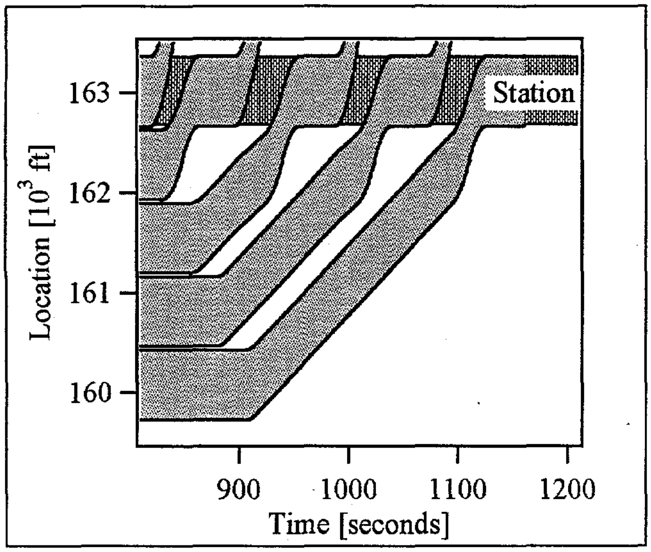

Figure 13. Enhanced control after backup

Coordinated Starts and Stops / Hills

It may be possible to reduce energy costs by more efficiently taking advantage of regenerated energy from braking. There are two obvious methods to achieve this: (1) coordinate the timing of trains arriving at and departing from stations, and (2) coordinate trains climbing up and down hills, so that energy is efficiently transferred from the braking (regenerating) train to the accelerating (powered) train.

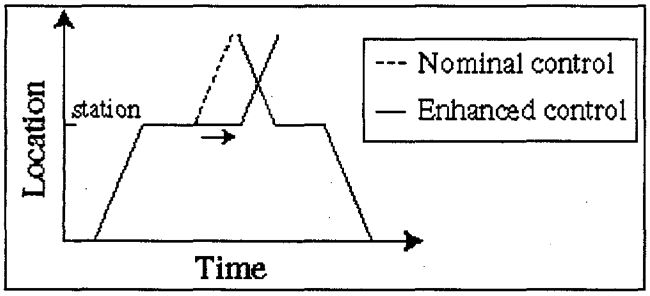

Figure 14. Coordinated starts and stops

A schematic representation of such an algorithm in action at a station is shown in Figure 14. In this case, the two trains are traveling in opposite directions through a station. The first train to stop in the station is held there by the algorithm to wait for the arrival of the opposing train. Thus, when the train is finally allowed to depart, it will use the energy available from the braking train rather than using power from a substation. Alternately, the approaching train could have been commanded to begin braking early to allow the departing train to leave on-time. Similar methods may be used to cause a train to climb uphill while another train traveling in the opposite direction is coming down the hill, braking to prevent acceleration.

We are not pursuing this algorithm at this time, because we do not believe the benefit of this control strategy would be sufficient to justify the extreme impacts on trip time. However, it would certainly be sufficient to justify developing a train schedule which encourages efficient energy transfer between trains wherever possible.

Coasting

BART trains typically travel at the maximum speed considered safe at any given moment, making it nearly impossible to reduce trip time any further. However, there is one exception to this rule: in some sections of track with long distances between stations, the speed limit is 80mph, but trains typically travel at only 70mph. This is done because trains traveling at 80mph and then braking from this speed can induce excessively large motor currents, which can lead to increased failure rates and maintenance costs. Coasting may provide a method to travel at over 70mph in these regions while avoiding the associated maintenance problems, and can not only reduce trip time, but can also conserve energy at the same time.

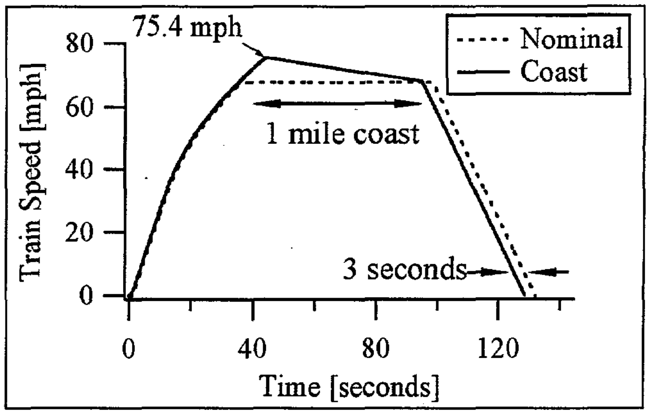

Figure 15. Coasting

A simulation of a one-mile long coasting trajectory on level track as compared to a nominal 70mph speed-maintaining trajectory is shown in Figure 15. The coasting trajectory was designed so that friction would slow the train to 70mph when it reached the location where it would begin braking. In this profile, the coasting portion of the trajectory corresponds to zero propulsion force, and therefore the speed drops due to drag and friction.

For the trajectories shown in the figure, coasting consumes approximately 10% less energy before braking begins than the nominal speed-maintaining trajectory, because the motors are more efficient under full propulsion than under low power while speed-maintaining. Meanwhile, coasting reduces the trip time by 3 seconds, because of the higher average velocity.

Some intelligence must be built in to any coasting algorithm, because the deceleration rate during the coasting trajectory will depend strongly on the wind velocity, grade, and perhaps track conditions. Airspeed in the simulation shown here is assumed to be zero, and a level grade is assumed, so this is a simplifying case. An algorithm would need to learn the wind speed, perhaps by recording the deceleration rate of successive coasting trains in a region, to adjust the speed before coasting as required.

Power-limited Acceleration

Power-limited acceleration is a technique which may be applied to reduce power consumption in some of the above algorithms. During acceleration, power consumption of a loaded ten car train rises to a maximum of approximately 8MW, and then remains fairly constant as a function of velocity. If the requested acceleration rate is reduced for a train accelerating from a stop, then this maximum power is reached at a higher velocity, and the train is slowed considerably. If, on the other hand, full acceleration is applied until some maximum desired power is reached, and then the requested acceleration rate is reduced to maintain a constant power demand, the resulting trajectory limits train power consumption while having a smaller impact on trip time.

Figure 16 demonstrates the difference between three acceleration modes: full acceleration at 2.8mph/s, half acceleration at 1.5mph/s, and power-limited acceleration with a maximum power demand of 5MW. In this case, the train accelerates to 34mph almost as quickly with power-limited acceleration as with full acceleration, but peak power demand is reduced. The half-acceleration trajectory never exceeds the 5MW limit either, but trip time is increased by 6 seconds. For acceleration to higher speeds, power demand does exceed 5 MW even with the half acceleration rate. The impacts of the three acceleration modes on trip time, and on the time duration when power demand exceeds 5MW, are shown in the table following the figure for acceleration to 34, 48, and 68 mph.

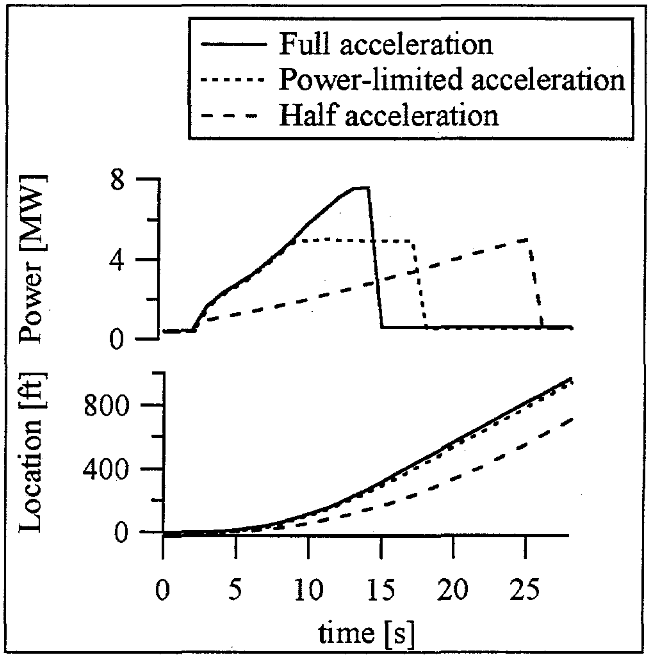

Figure 16. Acceleration to 34mph

Speed command	Accel mode	Trip time	time >5MW
68 mph:	nominal	-	31 sec
	half	+9 sec	27 sec
	5 MW	+5 sec	-
48 mph:	nominal	-	14 sec
	half	+7 sec	10 sec
	5 MW	+2 sec	-
34 mph:	nominal	-	6 sec
	half	+5 sec	-
	5 MW	+1 sec	-

Table 2. Comparison of acceleration modes

Power-limited acceleration may be used as a tool to regulate power consumption, for example in an algorithm to avoid low voltages. It is more effective for reducing peak power consumption than a constant reduction in acceleration rate, and has a smaller impact on trip time.

CONCLUSION

Development of a simulator of the train control and traction power systems at BART has been beneficial in assessing the benefits of conversion to a moving block control system from the present fixed block system. In addition, it has proven to represent an invaluable tool for developing and refining the control system before implementation, and for tracking down potential problems in the control system while it is still being designed.

The new Advanced Automatic Train Control system will allow not only more precise control of trains, but also coordination of the commands to multiple trains. Enhanced control algorithms will be incorporated into the system in order to reduce energy capital and operating costs, while simultaneously improving passenger comfort and reliability.

---

[1] Nishinaga, E., Evans, J.A., and Mayhew, G.L. (1994). Wireless Advanced Automatic Train Control. *IEEE Vehicular Technology Society News*, **41**, 13.

[2] Mayhew, G.L., Kivett, J.A., Himes, J.G., and Evans, J.A. (1994). Application of Radio Navigation Technology to Advanced Automatic Train Control. *IEEE Position, Location, and Navigation Symposium* (p. 217). Las Vegas, Nevada: IEEE.

[3] Lehrer, D.G. (1997). BART Train Simulator Documentation. Bay Area Rapid Transit District.

[4] Sa, T.J., and Gordon, S.P. (1997). AATC Train and Power Simulator (ATAPS) Documentation. Sandia National Laboratories.