

Appropriate Technology in Surface Scheduling

Peter Zelchenko

May, 2004

Originally prepared for presentation at Transport Chicago 2004

Easily the most pressing question for major metropolitan transit authorities today is how to deal with bus schedule reliability. Just a week ago, on May 12, 2004, the CTA board approved a test of GPS-based bus tracking technology, the newest component of the already partially deployed CAD/AVL complex — among other things, what is new is the part that transmits data back from the buses in real time.¹ The next step from there will be the “smart bus stop,” which in its turn receives bus tracking information and passes it to waiting riders. RTA is already moving to implement this.

Deals like the one above are, as we have seen, gradually and progressively brokered into the budget, primarily through an alliance between technology developers and boosters within the authority. In many cases, this costly technology fails to work as advertised or languishes due to improper use or unexpected and unbudgeted overheads. Once these technologies are approved, authorities and their management — as well as customers — are locked in for decades.

Many ingenious low-technology solutions are ignored, not because they have been proven to be unfit to solve the problem but because they are not attractively packaged or heavily promoted, as are the latest innovations. These low-tech solutions are not profitable to anyone; all they do is save money. Typically, a technocratic policy will advocate expensive solutions and leave no time for

ingenious thinking, sober leadership, human interaction, and best practices that lead to truly lasting and cost-effective solutions.

Technology is a mixed blessing

Case 1: Security

The best measures have usually been the simplest and cheapest. Transit security, though not a direct part of this discussion, provides ready examples because of the perennial potential to spend massive sums for negligible results, or tiny sums for significant results. Dummy camera

sconces and mylar lamination over the window glass has virtually eliminated graffiti on buses and trains. These two measures, arguably deceitful but tremendously effective, were far more cost efficient than, for example, New York’s “buff” car-cleaning plan of

the 1970’s, because they eliminated the problem of graffiti closer to the source, by addressing prevention. In many ways, the GPS-based management and information technologies of today come from the “buff” school of thinking, failing to address the problem of bus service with a mind to the law of parsimony.

As a more costly example, taken from general security, the monitors in the security booth in the Clark and Lake subway sit, uncalibrated, unfocused, dying, barely glowing, wasting electricity. This booth and others, each of

Many ingenious low-technology solutions are ignored because they are not attractively packaged or heavily promoted

Zelchenko is an engineer with nearly 30 years’ experience in multiple disciplines within computer science, industrial design, and graphic communications. He has written about and been interviewed on technology issues for the Chicago Tribune, the Chicago Sun-Times, Wired, the Boston Globe, and numerous other publications, and has commented at conferences and in published papers on appropriate technology. Zelchenko currently is doing research and development for on-demand book publishing technologies. He is also a lifetime user of Chicago transit. Peter Zelchenko, 1757 W. Augusta Blvd., Chicago, IL 60622, (312) 733-2473. E-mail pete@chinet.com, Web site <http://chinet.com/~pete/>.

which alone cost hundreds of thousands of dollars, were abandoned years ago. Each booth once housed a security person and still includes a number of monitors attached to cameras throughout the subway. Today, power is still running through this system, and several of the cameras are even still functioning, but there is nobody watching. Not long before this system was installed, the Chicago Police Department had a far more elaborate concept installed at a number of stations, connecting these stations by voice and camera to police communications at their headquarters at 11th and State. The pilot project alone cost millions and closed down after a short time. One of the complaints of these systems was that it is sometimes difficult to distinguish on remote camera between a mugging and an animated conversation between friends.

All of these systems eventually failed, much of the infrastructure has yet to be torn out, and people still get mugged at Clark and Lake. Tens of millions of dollars have been spent on these and similar projects over the years, and the reported benefits have been negligible.

Case 2: Fare media

I am continuing with a mention of farecards because it too is a case study of the ways in which technology can fail to fulfill predictions — of the subtle ways new fare media fail surface passengers by denying them information they formerly had and by failing to help solve the surface dwell-time problems it was claimed they would help solve.

Before the introduction of fare cards, cash, tokens, and paper-based transfers all gave the holder excellent information on the precise value of their assets. Even the

Flintstones-era see-through fareboxes, before electronic coin acceptors and digital fare media, were quite efficient. Riders used tokens or cash, or presented transfers. The rider always knew precisely how much money he or she had, and the farebox nearly eliminated driver embezzlement.

Even at this late hour, the CTA and other agencies still publicly consider new fare media an important contributor in their attempts to keep buses on time.² But, for bus bunching, these cards may do more harm than good.

Today, neither Chicago's magnetic nor its electronic fare media give riders any information on the value of their card; although there is ten times more circuitry in a single

Smart Card than there is in a digital watch, there is no display. Riders standing at a bus stop can never tell a \$100 card from a card with zero value — they must either go to a train station or wait to go through an arcane procedure with a bus operator — which regularly increases dwell time. In the meantime, it is no wonder that riders board the bus and they are a nickel short on their card. At the very moment in the system where the problem was to have been solved, riders are fishing for change or distracting the bus driver. Time and money slip away. It's worse than using cash, since with cash you know you are short at home or while standing at the bus stop. Note that this problem is not solved with Smart Cards unless a credit card is involved.

Earlier farecard system designs, such as the one used on BART and in Washington, D.C.'s Metro, at least printed the new value on the card. This acted both as a receipt and a visible indication of the value on the card. Increasing at bus terminals and train stations the installed base of machines that display the value only reinforces the point: Do what you will, portable fare media will be fundamentally flawed until there is some way to display the value on the card itself, a concept that is years away, if ever — not even on the radar.

“The CTA says it has also implemented many changes, including new technologies like fare cards...in an effort to speed up the boarding process and keep buses on time.”* But, for bus bunching, these cards may do more harm than good

***ABC-7 News, Feb. 16, 2004**

Finally, we still see problems with the farecard acceptors, such that it is not uncommon that an entire run will collect zero revenue because the farebox is out of order. I have heard of cases where a bus at a garage has been sent out, even though the acceptor is already known not to be operational, because no other bus was available and there was insufficient time to replace the farebox. The primary advantage, and to my

mind the only real advantage, of the Smart Card system is the elimination of media-path problems associated with accepting and scanning magnetic media and paper money. Smart Card allows for a scanner with no moving parts, which has been shown to translate to significantly less total device downtime. But this solution must be forced upon weary CTA riders by removing the bonuses from the magnetic card system.

One primary reason for moving away from cash and tokens was to move away from transfers. An important reason for moving away from transfers was to avoid the “Transfer Received” embezzlement scam, which was a human resources problem at the rapid transit stations and had nothing to do with the surface system. For buses, the

“Transfer Received” scam was not a concern because fares were always secured in the fareboxes. The touted advantage of magnetic cards in terms of buses was strictly to speed up the boarding process. But magnetic cards did not speed up the boarding process any more than tokens did, and in some ways it slowed it down and caused similar wholesale losses of revenue.

Case 3: Automated announcement

There are advantages and disadvantages to every technology that is introduced. Bringing the discussion closer to the surface scheduling problem, I mention the automated announcement systems now almost universal on buses and trains, in cities all over the world: The key advantages to these systems are that they allow the operator to focus fuller attention on operating the vehicle and they also offer a consistently clear and enunciated voice.

The Achilles’ heel of this technology is that it often is not smart enough:

Too often, on both the bus and train systems, the devices are miscalibrated to the stops or, if bus telemetry is lost, they are stuck on a certain stop announcement throughout the run.

Sometimes they are too smart.

One notable example is the programming for Chicago’s Clark and Lake rapid transit stop. During the rush period, the voice says, “This is Clark and Lake. Change to Orange, Green, Purple, and Brown Line trains at Clark and Lake.” Mention of the Purple Line, which only stops there during rush periods, is completely omitted at all other times. The system, we conclude, is intelligent enough to have a different set of messages at different times of the day.

At first glance, we think that this is a clever, perhaps ingenious, customization of the technology, like the first time we changed the background design on our computer’s desktop. But it is a mistake to use it so frivolously. Riders coming downtown, listening for some mention of the Purple Line — whether it be running or not — wonder what happened to it. Perhaps they are deliberately a couple of hours early, so that they may go shopping on State Street before heading to Evanston. When they hear where the Purple Line intersects theirs, they will detrain. But there is no mention of the Purple Line, so they stay on the train until they hear it mentioned — but it has mysteriously vanished from the sys-

tem, and they end up scratching their heads and staring at the map until they finally reach Chinatown and have to turn back. An hour lost through ostensibly clever customization. Sometimes features in technology can cause terrible confusion if used gratuitously or without much thought.

Snapshots of elegance

I mention South Shore Line’s fare payment receipts only as a convenient example close to home, to help us exercise the kinds of mental processes conducive to elegant and careful engineering thinking. When a fare is received, the conductor slides three pointed tabs into position on a small device. These tabs indicate dollars, dimes, and nickels of fare paid. When the receipt is torn off of the pad, it leaves three notches in the paper under where the tabs were, indicating the fare paid, and of course the stub shows the same information. At the end of the day,

adding up the pad stubs will indicate the conductor’s total fares received. This solution is perfectly portable and has about six moving parts.

This receipt system sits among the lowest-technology solutions available to transit, but it offers the greatest bang for

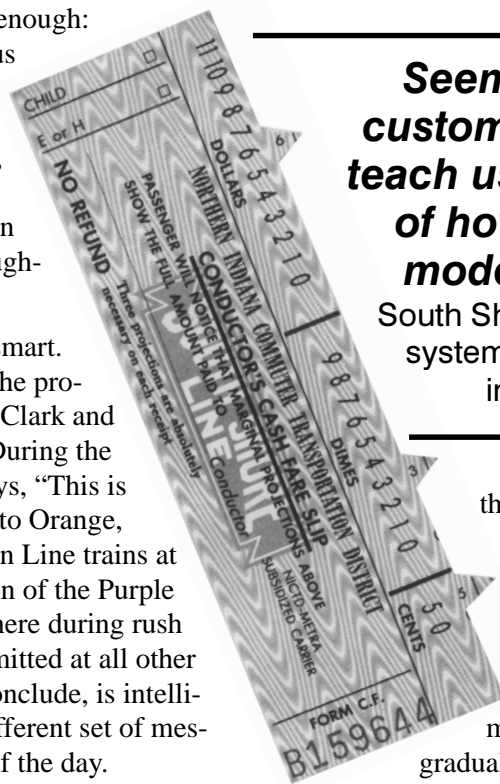
the buck. It is very low tech but very high value.

Seemingly archaic customer systems can teach us the philosophy of how to articulate modern technique

South Shore Line’s receipting system design is timeless in its elegance

A related engineering concept is the rolling-clock transfer devices still used by many bus systems worldwide, including some otherwise very high-tech ones like Seattle. A pad of transfers is held between two rollers attached to a clock motor. As time passes, the stack moves forward gradually along the rollers. An edge for tearing the transfers aligns with corresponding times printed along the edge of the transfer. As long as the time is correct and the transfers are properly in place, all an operator need do is tear the transfer off of the pad; the point at which it is torn indicates the time the transfer will expire. Unlike the last preprinted transfers CTA used — which required the printed date and route as well as punching of the time by an operator — no special printing or punching are required.

Although these concepts are tiny ones and perhaps are better suited for the annals of transit engineering history, they are still used extensively; this is because they exemplify the kind of inexpensive, evolutionary engineering



perfection that needs no further improvement and that has supreme longevity. Practically speaking, these concepts are, of course, inappropriate to our specific problems. But there is no reason that this general spirit of simple but ingenious engineering ingenuity cannot be expanded in scope, applied to the larger question of surface transportation problems.

Surface scheduling

The point that I hope was clearly made above is that technology includes many surprises, and sometimes for a large system the greatest surprise is that the technology was almost completely inappropriate. With this in mind, we finally arrive at the discussion of surface scheduling. The surface scheduling quandary in major metro areas is virtually synonymous with bus bunching, and in discussing the various high- and low-technology solutions I will begin by returning to fare media technology.

Although operators often give the complaint that handicapped riders cause delays which contribute to bus bunching, a much more frequent and less obvious delay is in waiting for regular passengers to board and pay their fares. It is not atypical for a single passenger to innocently impose a delay of 15 to 30 seconds while searching for their farecard or to add the correct change, or while negotiating with the driver, while a dozen passengers still stand outside.

One of the stated justifications for farecards was to speed up the boarding process and thereby reduce bus dwell times. As we now see, magnetic farecards did little or nothing to solve that problem, and it is unlikely that Smart Cards will speed up the boarding process enough to offer any statistically significant alleviation of bus bunching, if any at all. So, for practical purposes, officials and advocates should ignore claims of more efficient bus boarding with so-called improved fare systems. I anticipate that Smart Cards, even when deployed systemwide, will do virtually nothing measurable to alleviate bus bunching.

In terms of passengers lining up for fare payment and its contribution to bus bunching, the only known high technology that is reputed to be able to solve the problem is preboarding fare payment at the wayside, using some automated process. This has been attempted in some jurisdictions, but it is a very costly solution in terms of technology and human resources.

“Last on, first to pay” bus boarding

An astonishingly low-technology solution that can help solve this cause of bus bunching is a “Last on, first to pay” bus boarding convention. The cost may be as extravagant as a strip of yellow tape on the floor of every bus.

In a last-on, first-to-pay scenario, riders are instructed to board the bus immediately and line up, making room for other boarding riders, using what engineers call a last-in, first-out queue. The last rider to board pays first. The bus then is able to leave the stop immediately on boarding the last passenger, as there is absolutely no delay caused by fare handling, regardless of the fare medium used. It will be typical for a bus to save anywhere from 30 seconds to two minutes at many stops, and more at very busy stops. This single technique could yield the most significant reduction in bus bunching yet to be identified.

Dynamic routing

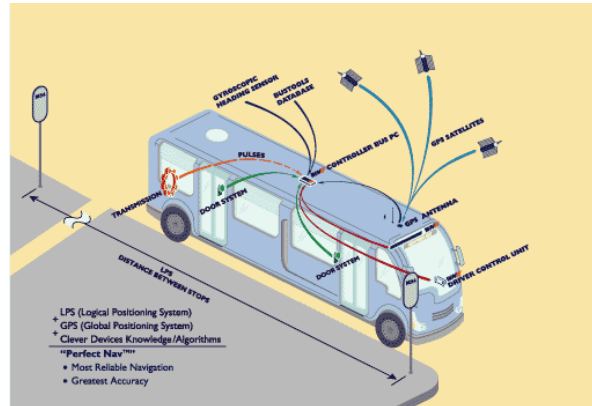
Much has been said about dynamically routing buses. There exists, for example, the option to change a run from regular service to express in certain circumstances when a bus is running far enough behind. More often used in rapid transit to alleviate serious delays, in bus practice it is rarely done because of its capacity to confuse and strand more riders than it rescues. However, it should still be considered as a tool in an arsenal of options for restoring proper spacing between buses. Unlike rapid transit express, an express bus can either eject all non-express riders at a given stop or allow riders already on the bus to alight at their designated stops but pass waiting riders during the switch to express.

Other low-tech dynamic routing options include rerouting along a less congested thoroughfare,

looping back, falling back or rushing, leapfrogging, and reassigning a bus, on the fly, to a different route entirely. Some of these are questionable as to their efficacy, but they should be considered in any complete service quality process.

Operator empowerment

It has become clear that a few route bunching problems cannot be solved by a standard arsenal of options and require special attention. But route management tools should be made available to all operators. A standard set of procedures should be clarified as to what options may



Clever Devices' SmartBus technology wants to make bus schedules obsolete

be employed under what circumstances. Operators should not only be given the freedom and encouragement to employ these options at their own discretion when justified and with proper training, but operators and teams of operators should also be rewarded when they use them properly and thereby contribute to improved service.

Operator-to-operator and operator-to-garage

Operators should also be encouraged to use the radio to communicate with leader and follower buses to control bus headway and ensure satisfactory service. An operator who knows his follower is only three blocks behind may not only feel compelled to pick up his pace, but should also feel less concerned about passing up a waiting customer to keep on schedule. A relatively low-tech way to reassure waiting customers in this circumstance would be to customize the bus's front forehead display to alternate the route number with "RUNNING LATE / NEXT BUS 4 MINS" or something to that effect.

CTA has attempted this kind of decentralization of management a number of times in the past. Most recently, in 2000 it implemented an Operator Empowerment Program to "[allow] bus operators to make decisions about service restoration techniques that previously could be made only by Control Center personnel or a supervisor on the scene."³ They offered operators a number of options, including "spacing back to leave more distance between buses; leaving the terminal early to fill a gap in service, or detouring on the spot to bypass a blocked street." Even the limited choices granted to operators during the pilot period yielded reductions in bus bunching of between 1.5 and 3.5 percent, depending on route length and complexity. Further empowering operators with an entire battery of options, based on a serious analysis of the possibilities, should increase these numbers significantly.

37 SEDGWICK
 Monday thru Friday
 Southbound
 Northbound

For more information
 Travel Information Center
 312.463.6200

am light face
 pm bold face
 No Saturday, Sunday or holiday Service

Emphasis on online schedule availability ignores the fact that there is an inverse correlation between the surface customer demographic and its ability to access the Internet.

Virtually all of the benefit of this type of thinking will come from "no-tech" sources, primarily from robust training and the reliance on the operator's own mental faculties. The only technology that need be employed to enhance operator empowerment would be the bus's onboard radio, whose functionality has not changed radically since first implemented on buses. This may be contrasted with the centrist and CAD/AVL-oriented approach of bus management, which will cost in the high five digits per bus and include annual maintenance costs, including significant human resources. These huge hidden costs have not been clearly listed in any planning budget that I have seen for these new technologies.

Appropriate scheduling management

Most important to a treatment of bus bunching is a steadfast commitment to following the advertised schedule. Drivers should be rewarded for maintaining their schedules, but this reward system should be specially weighted based on the difficulty of the route. The best drivers should be promoted to the most difficult routes and compensated accordingly.

More importantly, if a route is regularly not keeping to the public schedule, and efforts to adhere to the schedule continue to return double-digit inefficiencies on many routes, two choices remain: either remove the schedule entirely or revise the schedule. (Adhering to an existing schedule can primarily be effected by improving travel time, which discussion see below under "Traffic management.")

Remove the schedule

While in Chicago's example the CTA publishes nominal schedules for all routes, it has never made the schedules very visible to riders. Perhaps nothing is more frustrating to a rider than to know that a bus is intended to keep a schedule but to have the schedule nonchalantly hidden from view.

The vagueness at CTA bus stops — from the mysterious "Day through early evening hours" to the ominous "No Owl Service" to certain State Street stops which have no information at all, let alone such deliberately vague lan-

guage — may save hundreds of thousands of dollars annually in printing and distribution costs but is a terrible injustice to riders.

The response that schedules are available online raises the question of why they are not available offline, particularly at the bus stops and on buses, where they are needed. It ignores the fact that there is an inverse correlation between the surface customer demographic and its ability to access the Internet to obtain schedule information. This is not a good place to save money, and it does not contribute to rider confidence. In fact, even at CTA headquarters itself, typically so many schedules are missing from the display that it would be a waste of time for most riders to come in to find one for their desired route.

Already, when riders today call the CTA, they are told only that they can get schedules online. This is already very inconvenient or impossible for many riders. They will mail a rider a schedule and system map on request, but schedules are not available on buses and connecting rapid transit stations where they belong. Tomorrow, printed maps may vanish completely:

“The [CAD/AVL] information would be posted on the Internet, accessible by computer, cell phone or personal digital assistant. It would also be fed to the RTA’s travel information phone center and transmitted on electronic message boards and kiosks at *some* high-volume bus stops and train stations, although wide-scale implementation is still *at least several years away*. ”⁴ (Emphasis mine.)

Clearly, even if these devices provide schedule information, it is so cost-prohibitive to put them at every bus stop that the only alternative is printed schedules. But the kind of talk we read above suggests that we are getting further away from printed information at bus stops and on buses, not closer. The next step is to declare it obsolete and to be advised that we should consult Candace Bergen about buying a cell phone.

The proper place for schedules is where they are needed most. Detailed schedules should at minimum be displayed broadsheet-style on every bus and at every bus stop. Copies should always be available to customers in these same places, as well as online. Every operator, as part of the procedure, should insert the appropriate route map and schedule into a display on the bus, and check for copies in the bins. Ideally, maps and schedules should be available in establishments along every route.

All of this is done in smaller jurisdictions and one would assume that such a practice would scale economically in such a large city. If the counterargument is presented that Chicago’s system cannot afford to print so many schedules that change so frequently and are so unreliable, they should still post up-to-date schedules on every bus and at every bus stop. Although this was attempted at a few bus stops, it has been abandoned. Riders should be encouraged to know how late a bus is and should be encouraged to notify the CTA of any problems with service. This information should not be hidden from riders, because the result is not only that riders are frustrated, but that management gets even less feedback than they currently do. Another missing piece of the puzzle: a rider satisfaction suggestion box on every bus.

Revise the schedule: A counterintuitive strategy

Probably the most common excuse about the poor dissemination of this information is the claim that the kind of public information offered by published schedules, while effective in small-system jurisdictions for their lower-frequency routes, is almost irrelevant in a major metropolitan system in which buses are scheduled to arrive every 7 to 12 minutes. Furthermore, it may be whispered, it is a disservice to promote a schedule that cannot be reliably adhered to a large percentage of the time because of such a congested system.

These may be valid explanations, but it may raise the question: Why are we running so many buses if they get bunched so often, even off peak? Admittedly, rush-hour resource allocation makes bunching virtually unavoidable in many cases, and I submit that no technology can solve this. Several buses must come in tight succession to accommodate a surge in ridership to and from the business district and large schools and other institutions. However, most off-peak bus bunching is theoretically preventable, and many other deviations from the schedule are definitely preventable.

So, again the question: Apart from rush periods, why run so many buses? It is because of the institu-

tionalized expectation among city riders that buses on most routes are intended to arrive “every few minutes” all day, and the concomitant need to increase runs because bunching was causing unacceptable waits of up to an hour and sometimes longer. We did not remove the cause, we only treated the symptoms by increasing the dose; costs rose astronomically, other services have been cannibalized, and the problem is still not solved.

It may seem heretical, but it is possible to alleviate the bus bunching problem by actually *thinning out* off-peak

It may seem heretical, but it is possible to alleviate the bus bunching problem by actually thinning out off-peak service systemwide

service systemwide, even on major routes. But customer satisfaction will only be improved if this is done in combination with vastly improved access to schedule information (particularly printed schedules) and sticking religiously to the more manageable schedule. Off-peak riders, if they can rely on the fact that a bus will arrive at their stop at a particular minute most of the time, will be able to plan their travel with lower pressure and almost no time waiting at the stops, much more in the style of surface transportation seen in smaller communities.

The difference in frustration levels between riders of small and large surface systems relates to how long they are forced to wait for an unknown quantity. Small-system riders saunter out to a stop at 3:11 p.m. every day and board their bus at 3:14. Large-system riders walk out to a stop at 3:11 p.m. and could board at 3:14 or 4:00, or hail a cab if they're relying on obsolete or obscure signage or no signage at all. Removing the unknowns will eliminate the waiting time.

Many riders, if not most, may prefer to see half or even one-third the frequency of buses on their route if they were promised precisely when a bus would arrive and could rely on it.

The theory that headway problems are increased by decreases in design headway has been considered in the past and has more recently been confirmed by computer simulation.⁵ But this research has typically been used to justify high-technology control solutions rather than considering reductions in run frequency.

Key to a reductionist approach is the commitment (detailed above in "Remove the schedule") to keeping schedules accurate and available. It will not do to have bus schedules only available at CTA headquarters (and there only half of them available at any given time). Schedules must be ubiquitous; that means that it is far more important that they be available on paper than on the Internet, more important that they be posted religiously at every bus stop and on every bus than on cell phone displays.

As a kind of "reverse technology" approach, this concept can lead to enormous efficiencies and cost savings by greatly reducing the needed number of runs on even the city's busiest routes. Here we are presented with a huge opportunity for large surface systems: With the savings realized by this fundamental change in philosophy,

round-the-clock service can be extended. Short routes can be extended. New shuttle services can be introduced. All of these changes can be implemented by reassigning existing operators. A system could exploit the Internet and telephone in a sensible way, not to tell riders how late a bus is going to be, but to allow for the introduction of customizable small-shuttle service, almost as one might for library books: ask for a shuttle to pass any given point for the next day, two weeks, a month, at any given time. This is technology placed appropriately.

Twenty years ago, this was the vision for the surface transportation systems of the future, but we have become too caught up in the excitement of GPS, putting technology where it doesn't belong. GPS is not inherently a management tool, and it is potentially a mismanagement tool. Riders don't need to know where buses are when they aren't where they are supposed to be, they need buses to be where they are supposed to be. With GPS, the dog walks the master.

GPS and wayside AVL

The technocratic manager will argue that, in order to maintain control, management must have as much information at its fingertips as possible, at an instant's notice, regardless of the cost and the amount of extraneous data generated. Lacking innovations to solve bus bunching, management has been left hanging with their proverbial flies open. Like the other half of a zipper, GPS vendors

Many riders, if not most, may prefer to see half or even one-third the frequency of buses on their route if they were promised precisely when a bus would arrive and could rely on it

The lone senior citizen, standing freezing at a desolate street corner on a January evening, will not be consoled by the little flashing blip, stuck on the far left side of the screen

have rapidly evolved to close up the embarrassment, by cozying up to the technocrats in transportation in order to accommodate a mode of thinking more attuned to that of the centralist management theories of the mid-20th century. But one must not confound technology and innovation, and one must not assume that every idea, strictly by virtue of its novelty or diligence, is a valid solution. What purpose does it serve to know precisely where every bus currently is in the system every minute of the day?

Certainly, the statistics gained and the possibilities for modeling this data are potentially very valuable, if it is going to solve the problems. But management already knows very well where the problems are; they already have a massive quantity of data from a large number of sources. GPS technology will merely give them updates cheaply and cleanly, and at the push of a button. In itself, it solves no problems.


And now, transit authorities all over the world are poised to take the most costly plunge into this fallacy: putting wayside arrival systems at every bus stop in every city. This is the culmination of the technocratic dream, since for managers it represents the greatest possible information dissemination to date, and for vendors the greatest possible technological expansion. It is fabulously expensive. It appears as a logical, natural, almost organic appendage to all of the solutions attached to the Swiss army knife that is GPS.

What information is being given with wayside displays? Aside from the gratuitous weather report, time, and other window-dressing, the centerpiece is that riders will be told when the next bus is *expected to arrive*, whether it is on time or not. With all of the batteries and printed circuits, with all of the radio telemetry and signal modulation, wayside systems cannot change the fact that the bus should have arrived half an hour ago. The lone senior citizen, late for her weekly event, standing freezing at a desolate street corner on a January evening, will not be consoled by the little flashing blip, stuck on the far left side of the screen. She has been standing there for decades; to her, the display is cold comfort and offers nothing she didn't already know — that her bus is and always will be very late.

Though technocratic thinking is innocent in its shortsightedness, wayside arrival monitors do not solve fundamental problems of bus bunching and they allocate enormous resources that could be better used actually to solve the problem.

Traffic management

The hottest new technology for surface traffic enhancement is radio-modulated traffic signal prioritization for buses. This is not quite as sexy a technology as those listed above, but it is far from cheap. Furthermore, this solution gives not only buses priority through the intersection, but also the many cars that may be causing the congestion in the first place.



This solution yields better results on suburban routes where lights are longer and traffic less congested. Lower, more "unplugged" technologies may serve better in the city and leave some money to reverse service cutbacks.

One of the most significant of these, as evidenced by the fact that it has already been implemented on a broad basis, is far-side stop relocation, in which bus stops are put on the other side of the intersection — after the light instead of before it. This allows the bus to move through the intersection on the green, when the light is green. Even including new pouring of concrete and moving signage, dealing with aldermen, transportation, and revenue departments, and all of the rest of the red tape, in the end this is such a relatively simple, intuitive, and inexpensive matter that it is difficult to pass up.

Implementing this solution alone will not eliminate bus bunching. But this in conjunction with a number of other practical low-tech solutions could reduce bus bunching to such a degree that CAD/AVL becomes irrelevant on the basis of more than simply its own limitations.

In this vein of thinking, bus priority lanes should always be reviewed and revisited. This class of idea was used more often in downtown Chicago before the 1970s, and it disintegrated as traffic congestion called for more lanes for automobiles. A concept that has not been considered is, along appropriate routes, a center lane usable only by buses and turning vehicles in both directions. Added to this can be a painted region through the car lanes which cannot be occupied by cars when buses are present (see image), allowing buses to move immediately to the bus stops instead of having to wait for cars to clear. On many streets with relatively narrow rights of way, implementation of this design including the third-lane option would have an added traffic-calming effect.

One idea, involving thinking that is quite out of the box, is to mount Mars lights on buses and to enforce law that gives buses priority over cars. This would be legally justified based on a high-occupancy superior right of way for the buses. As with emergency vehicles, cars would be forced to pull aside to allow buses to pass.

New York's Transit Alternatives group, in its push for Bus Rapid Transit practices,⁶ recommends

some additional possibilities along this line: "bus bulbs, which bring the sidewalk out one lane so buses do not have to maneuver into and out of bus stops, longer bus stops to eliminate delays as buses wait to enter the stop,

and bus lanes with raised lane dividers or other physical means to discourage or prevent other vehicles from violating bus lanes.”

This class of design engineering is full of potential, and unexpected, heretofore unconsidered ideas are simply waiting to be discovered and tested.

Appropriate computer-aided “just in time” route analysis

A classic contribution to bus bunching occurs when the turnaround point for a route is in a congested area, such as in or near the central business district. As an already compromised schedule stacks late buses up near the route’s turnaround point for the return run, buses are so far off of their schedules that the printed time is totally meaningless on the return trip. In fact, a clock-time-based schedule will be meaningful only to the extent that the operator can mentally keep track of the number of minutes he or she has deviated from perfection and make mental adjustments. Any amount of deviation requires unnecessary processing.

There are at least two solutions to this problem using technology appropriately. The first is to abandon schedules with specific printed times for the operators, at least on routes that regularly experience bus bunching. The times listed on these schedules will be meaningless to a driver on a route that is even a few minutes late in starting. To give more meaningful data, one of two things can be done. A zero-cost solution would be to print the provided operator run schedule with a T-plus minute position to the side of the actual time, and perhaps also the time in minutes between each milestone. If the run is off time, drivers will use a stopwatch calibrated not to the clock time but to the last known milestone and work from there. Alternately, operations management could couple the existing onboard passenger information data and a stopwatch function built into the onboard computer’s firmware to give drivers an easier way to gauge spacing on the return trip. But this would necessitate that the computer be working and properly calibrated, which problem is discussed elsewhere in this paper.

In either case, every driver should coordinate with his or her follower by radio to establish their new starting time.

The goal should be to reestablish the proper headway if one run or a large number of runs are off schedule. This is not a complicated strategy.

Even better than the basic T-plus column recommended above would be a printed T-plus minute timetable for an entire day, based on each run’s actual start time. The list would be a duplicate of the page from a garage’s master scheduling data but with T-plus information emphasized. An operator would observe his or her time out from the starting point and follow that time’s T-plus data against

the stopwatch. The advantage to this, particularly across the transitions between rush and non-rush periods, is that it takes into account nonlinear changes in the schedule throughout the day; buses thus adapt more precisely to changes in time factors when starting out later than expected.

This class of solutions, which increases operator autonomy in the management and coordination of scheduling, can be extended if sufficient information is

provided. Taking the example of a route that is already compromised in rush period and is backed up at the turn by more than the number of needed buses, one driver can be instructed (or can himself or in consultation with other drivers elect) to switch the bus to a different route which could use the additional resources, based on a book (or, yes, even an onboard database) of all routes which start at or near the bus’s current location and a plan of action. While this is already done occasionally, proper management exploitation of this technique could significantly reduce the resource waste that is inherent in bus bunching.

In modern management parlance, these practices would be termed “just in time” procedures. But many managers assume incorrectly that just-in-time management requires high technology at every point in the system. Readers may question whether these are low-tech solutions if they involve computers at all. I tend to agree; neither the single-route T-plus solution nor the extended master timetable rerouting solution requires the onboard computers, since the same purposes can be accomplished, and perhaps more reliably, with a \$5 stopwatch, using printed charts generated by a very modest computer and database.



Clever Devices (left), NextBus (below), and Orbital-TMS (similar to NextBus) are pushing hard to sell wayside solutions.



However, while this critic believes that, given two solutions whose desired results are identical, the lower-technology solution will usually be more reliable and often more convenient to use from the perspective of design, it should be admitted that if a technology is already paid for and available, it would be a waste not to exploit it to its fullest extent. It must be said here, however, that both of these solutions employ some form of technology; the onboard computer solution simply employs far more technology, deployed at far more positions, to yield the same result, and thus is in many ways less efficient. The question at hand in this entire paper to address what technology is *appropriate* for a given operation.

It should by now be clear that it is not absolutely necessary that techniques like these employ onboard technologies, even if they are available. The best solution might be to train operators in both, so that they might understand not only how to employ the technique, but the idea behind it, and can learn to improvise intelligently if necessary. With gratuitous high-technology solutions up to the eyeballs, we not only insult the intelligence of operators, we also hand them a license to stop thinking.

Getting the rider to buy in

Studies, including one by the Federal Transit Administration, are showing what we might expect — that, for riding customers, “the preferred information delivery systems tend to shift from static scheduling information to real-time vehicle arrival/departure time and location updates.”⁷ The FTA study ignored, or may not have clarified to subjects, the fundamental difference between “static” and “real time” information. The two concepts rely on core differences in transit policy and philosophy. The first articulates a system based on a highly predictable schedule for which any deviation would be considered an exception; the second is based on the notion that there is no real schedule that need be committed to. It is as if, absent solutions for adhering to a static schedule, we were attempting to force a preference for the second over the first and package it as the wave of the future.

These solutions are being pushed for the ADA primarily because that is where the money and voice are. But pro-

ponents for these groups are also being shortsighted for the same reasons: these solutions do nothing to solve the underlying problems.

A priesthood and the hidden costs

I have not seen an implementation of a technology-based system that has not involved significant additional charges amounting to several times the back-of-the-envelope costs for implementation. One rough estimate puts the true total cost of a technology at the original sticker price as not a single fixed cost, but as a multiplier to be incurred every year, plus extra for the first and last years of service life (for installation, network, training, unanticipated parts and labor, removal, etc.). For example, a \$1,000 computer system will cost an additional \$1,000 in the first year for wiring, installation, shipping, taxes, training, and additional materials, and an additional \$1,000 every year for ongoing costs, including power and network, and including special human resources and management, telephone, and shipping costs of upgrades, maintenance, and leasing during failure periods. It will also cost \$1,000 to remove and properly dispose of the equipment. So, for a device whose original cost is \$1,000, in this model, the total cost over a five-year lifetime is closer to \$7,000.

These costs include the added expenses for training and for an expanded priesthood of highly paid, usually degreed, staff to manage the technology. It includes the costs to upgrade the equipment and to replace it, and the real and cash-equivalent losses incurred over time when a portion fails. It includes the physical and environmental cost of removal of the system when it either actually reaches end of life or is abandoned, either in favor of an upgrade, because it failed to meet expectations, or because of unanticipated budget cuts, including for staffing to

manage and maintain the technology. Those technological implementations that have not been given major annual infusions for upkeep and upgrade have gradually, through disuse, sunken into the sands of time.

A good example from Chicago’s transit history of these costs can be seen in the subway security systems mentioned in the introduction.

Bus tracking is “a necessary step in the evolution of our service,” said CTA chairman Carole Brown. “We simply must solve the problem of bus-bunching and provide our customers with service they can count on.”* But the technology in itself does nothing for either of these things

***Jon Hilkevitch, “CTA Expected to Approve Bus-Tracking Test,” Chicago Tribune, May 10, 2004**

Conclusion

At some point in civilization's early period, potters and Bronze-age craftspeople began to put handles on cups. This allowed diners to hold hot drinks without burning their fingers, but it presented a new potential point at which every cup could break. Every new circuit, every new special instruction, is the introduction of a new potential point of failure in an already deeply complex transit system. This grows exponentially with higher and higher technology. The total reliability of a system is normally inversely related to its complexity, and the law of parsimony must always apply when considering transit improvements.

GPS and CAD/AVL will not of themselves solve bus bunching any better than will a sharp operator or dispatcher ever could, and they have already created a costly new culture of urban travel planning that allows at least the technologically privileged to excuse unreliable service because they, at least, know exactly how much longer they need to wait. For management, this technology will be only as effective as those who master it, and in the long run it is totally disempowering of operators.

Every technological solution is an indirect move sideways, distracting attention from the source of the problem. If the system can be improved more directly, it should be. Improving the human element, by truly empowering operators and riders with information, by giving operators incentives for rewarding riders with reliable service, by creating a partnership among operators and riders, is a difficult route, but it is the most direct route to making the system reliable. Simple engineering — implementing solutions as simple as spreading new concrete, paint, paper and ink, and moving bus stops a few feet in one direction or another — is in the long run far more effective and less costly than telling our cell phones how late the next bus is going to be.

Simple engineering is also more broadly beneficial. It is true that the general attitude that the world's fast-paced schedule will be linked to our cell phones has merit for some. A group of five executives who all are carrying cell phones can push forward a meeting's start time by an hour with just a few minutes of phone tag or a few text messages. But most people do not live like that, particularly those who most often depend on bus transit. They have had meetings to attend as well, ones whose conveners have often been less sympathetic to the misfortunes of a repeatedly late bus.

Some transit officials will say that they have tried a number of low-tech solutions at one point or another. Some are being implemented now, but for some reason officials don't imagine that the problem can theoretically be solved with a package of low-tech solutions. The most tedious thing to do in a high-tech world is to study and implement the low-tech solutions. They are not sexy and they often don't *appear* to be benefiting anyone directly. The least tedious thing to implement is usually the most expensive — the thing that is being promoted by outside boosters and for which these boosters promise a very

smooth transition. These solutions, they claim, have worked for other authorities. Sometimes the vendor incentivizes the deal, making it look like a "bargain."

The astute reader will notice that, except for the right-of-way discussions, one common thread within the low-tech position is to

encourage greater communication, particularly operator-operator, operator-rider, and authority-rider. This is in contrast to the more bureaucratic, higher-tech mode of management-to-operator.

What these techniques do is eliminate the middleman. In the case of conventional management, the middleman is the manager who must be consulted for every proposed change in plans. In high-technology management, the middlemen include not only management but also vendors and the authority's own expanded technical staff.

"Last on, first to pay" boarding, far-side stops, predictable scheduling, robust availability of printed passenger notification, basic right-of-way priorities, and interoperator communication. Other than two-way FM radio, which has been available on buses since the 1960s, none of these solutions involves one new copper wire or transistor, and by extension these involve no annual vendor maintenance fees, no downtime, no priesthood of technicians, and no upgrades. This kind of thinking can virtually eliminate the need for high-technology stop-gap measures. What is needed, aside from thorough examination of every option available, is a carefully done computer simulation of each of these technologies in combination, to see what set of solutions promises the best total system improvement. This, at least, is appropriate use of high technology.

Simple engineering is in the long run far more effective and less costly than telling our cell phones how late the next bus is going to be

Notes

1. Jon Hilkevitch, "CTA Expected to Approve Bus-Tracking Test," Chicago Tribune, May 10, 2004.
2. Kevin Roy, "The Waiting Game," ABC-7 News, Feb. 16, 2004: "The CTA says it has also implemented many changes, including new technologies like fare cards and talking buses that announce the routes in an effort to speed up the boarding process and keep buses on time."
3. Chicago Transit Authority, "CTA Expands Program to Reduce Bus Bunching," press release, October 4, 2000.
4. Hilkevitch, loc. cit.
5. See, for example, P. Chandrasekar, Ruey Long Cheu, and Hoong Chor Chin, "Simulation Evaluation of Route-Based Control of Bus Operations," Journal of Transportation Engineering, Vol. 128, No. 6, November/December 2002, pp. 519-527.
6. Transportation Alternatives (NYC), "Bus Rapid Transit for New York City: A Summary," <http://www.transalt.org/info/brt2.html>. "BRT features showing the most promise for implementation in New York City include: More frequent service where needed. Bus bulbs, which bring the sidewalk out one lane so buses do not have to maneuver into and out of bus stops. Longer bus stops to eliminate delays as buses wait to enter the stop. Bus lanes with raised lane dividers or other physical means to discourage or prevent other vehicles from violating bus lanes. Low-floor buses that can speed boarding and exiting and encourage riders to exit through the rear door. Pre-boarding fare payment at selected high-volume boarding times/locations to reduce dwell time at bus stops. Bus traffic signal priority to help late-arriving buses catch up to schedule. Real-time management of buses to achieve even spacing between buses. BRT features applied to M15 limited stop buses on First/Second Avenue in Manhattan could dramatically improve bus service. Note that many of these are low-technology solutions, and "real-time management" does not presuppose GPS. "BRT has been applied successfully in major cities including Los Angeles and Vancouver, British Columbia as well as cities in South America, Europe and Australia. BRT has produced 15-40% increases in bus speeds and 15-150% increases in ridership. ...One alternative, using dual bus lanes, low floor buses, raised lane dividers and pre-boarding fare payment during rush hour at six locations, would reduce bus travel times by 21-27% compared to the current limited stop service and reduce the variability of travel time by 38%. Bus riders would save 9-17 minutes for a trip from 125 Street to Houston." Note that most of these are low-technology alternatives.
7. Federal Transit Administration, "Preferences for Transit Traveler Information: Overview of Research Objectives and Findings," Customer Preferences for Transit ATIS Research Report.