PERFORMANCE AUDIT OF TRANSIT

TECHNICAL REPORT B:
SERVICE DEVELOPMENT

King County

Presented to
the Metropolitan King County Council
Government Accountability and Oversight Committee
by the
County Auditor's Office

Cheryle A. Broom, King County Auditor
Ron Perry, Deputy King County Auditor
Kymber Waltmunson, Principal Management Auditor
Booz Allen Hamilton, Consultant
Nelson Nygaard, Consultant
Courval Scheduling, Consultant
MicroKinetics, Consultant

Report No. 2009-01B
September 15, 2009

Auditor’s Office Mission

Through objective and independent audits and services, we promote and improve performance, accountability, and transparency in King County government.

Auditor’s Office Vision

Our work is of the highest quality and integrity resulting in significant improvements in accountability, performance, and efficiency in county government, and it promotes public trust.

❖

The King County Auditor's Office was created in 1969 by the King County Home Rule Charter as an independent agency within the legislative branch of county government. Under the provisions of the charter, the County Auditor is appointed by the Metropolitan King County Council. The King County Code contains policies and administrative rules for the Auditor's Office.

The King County Auditor's Office provides oversight of county government through independent audits and other studies regarding the performance and efficiency of agencies and programs, compliance with mandates, and integrity of financial management systems. The office reports the results of each audit or study to the Metropolitan King County Council.

The King County Auditor’s Office performs its work in accordance with applicable Government Auditing Standards.

❖

Audit and study reports are available on our Web site (www.kingcounty.gov/operations/auditor.aspx) in two formats: entire reports in PDF format (1999 to present) and report summaries (1992 to present). Copies of reports can also be requested by mail at 516 Third Avenue, Rm. W-1033, Seattle, WA 98104, or by phone at 206-296-1655.

Alternative Formats Available Upon Request

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Service Development Involves Balancing Customer Service, Operating Cost, and Operator Working Conditions

Scheduling service for a transit system involves striking a balance among three elements: excellent service to customers, the lowest possible operating cost, and providing satisfactory operator working conditions. In many cases, Transit's service choices have put more emphasis on frequent, timely service to transit customers and favorable working conditions for operators. These choices have come at an added cost. This cost and opportunities to increase the system's efficiencies are described in the following technical report. We recommend actions that, if implemented, could result in annual savings of $16 to $23 million annually following the implementation period.

This report consists of six chapters. The Introduction will give background information on Transit's service development process and an overall summary of our findings and recommendations as they relate to Transit bus and trolley service. The second chapter reviews opportunities to improve the strategic approach to service development, and the third chapter describes an efficiency-focused review of the bus system as a whole. The next chapters focus on building cost-saving approaches into the three phases of developing a service schedule: service trip definition, blocking, and runcutting. The final chapter looks at the software that Transit uses to develop its service and opportunities for utilizing it more effectively to ensure the most cost-effective bus service.

1 INTRODUCTION

Summary

This chapter provides background on Transit's schedule development process and the software used by Transit to facilitate this process. It also describes the objectives and methodology used in analyzing Transit's service. The chapter concludes with a summary of the findings and recommendations related to Transit service and an accounting of the estimated savings that could result from implementation of these recommendations.

Service Development

Transit Updates Routes and Schedules Three Times a Year

Transit updates its routes three times a year – in February, June, and September. Transit also takes this opportunity to reallocate buses and operators. There are three phases of the process that results in the development of a new schedule:[1]

1. Service trip definition- When Transit planners and schedulers define service trips, they identify the routes that buses should take, how long it takes each bus to complete its route, how frequently buses should run down each route, and key connection or transfer points.
2. Blocking- Blocking activities take the information developed in the first phase of scheduling, service trip definition, and assign vehicles to each service trip to form vehicle "blocks."
3. Runcutting- Finally, schedulers take each block and assign them to a "piece of work" that will be assigned to an operator. Operators then go through a "pick" process in which operators choose pieces of work based on seniority to

[1] The process of developing and implementing a new schedule is sometimes called a "shakeup" or "pick."

determine which routes they will drive until the next schedule is developed. This process is heavily controlled by the labor agreement.

These three steps are generally sequential but also inform each other iteratively. Most of the activities take place in Transit's Service Development Section, but the Vehicle Maintenance Section provides information about the coaches (vehicles) in Transit's fleet, and the Operations Section adds information about the characteristics of the current operator work force.

Scheduling Software

Bus Scheduling Is Highly Technical and Data Driven

Transit owns a scheduling software system, HASTUS, which is used in several other large transit agencies in North America. HASTUS assists Transit personnel in the highly technical and data driven process of schedule development. It is a complex scheduling package with sophisticated algorithms that have the ability to create very efficient transit schedules. Transit must program the software with appropriate information to ensure that it produces usable results. For example, HASTUS must "know" about the conditions inherent in the current labor agreement, available fleet, local geography, and other information about Transit's unique environment. Because of its complexity, a very high level of expertise with the software is required to produce the most effective results.

HASTUS is a large database with multiple modules which are designed to work together to address different tasks and problems within the scheduling process. In this report we discuss just a few of these modules specifically:

- The *ATP* module assists in building appropriate service trip definition by analyzing data from Transit's automatic vehicle location (AVL) system and current schedules.

- *Minbus* is a module with advanced optimization tools for vehicle scheduling during the blocking process.
- *CrewOpt* creates and optimizes operator assignments during the runcutting process.

Objectives and Methodology

The entire Transit audit spanned multiple areas of work, including Transit's service design practices, financial and capital planning, technology and information management, vehicle maintenance, operator and transit police staffing, and Paratransit. The objectives of this portion of the Transit audit were to examine Transit's bus service in relation to the general efficiency and effectiveness of scheduling.

To achieve this objective, the office and its consultants:

- Interviewed Transit leadership, management, and service development section line staff
- Interviewed GIRO staff, the vendor of the HASTUS software
- Observed meetings held during the scheduling process
- Surveyed relevant industry literature and best practices
- Reviewed Transit documents and agreements
- Analyzed Transit data provided by the software vendor, GIRO
- Performed analysis of Transit data including data from
    - A copy of Transit's HASTUS database, system files, and historic information
    - Automatic Vehicle Location (AVL) data
    - Automatic Passenger Counter (APC) data
- Developed scheduling models using the HASTUS software for sample routes and bases

Summary of Findings

Transit schedulers have demonstrated a thoughtful approach to solving scheduling problems and work to develop schedules that

Bus Service Could Be More Efficient

can be achieved during typical operating conditions while providing a cushion in case buses are running late. Nevertheless, we found opportunities for improving efficiencies without necessarily cutting service. For example, current scheduling approaches have resulted in Transit spending more time resources than are required to maintain schedule reliability. Even with the current scheduling practices there are opportunities for service provision to improve.

There are opportunities to enhance the strategic approach to Transit's service. Transit does not yet utilize performance metrics that monitor operating cost efficiency over time. This limits their ability to set performance targets and monitor progress toward ensuring that resources such as vehicles and operators are allocated to achieve lowest possible costs within the context of overall service delivery objectives. In addition, Transit does not have specific standards or guidelines that establish a framework for the trade-offs between efficient operations and other scheduling objectives (such as on time performance or passenger crowding) to direct the systematic use of Transit data in building more cost-efficient schedules.

Although Transit schedulers use many effective methods for building schedules and assigning staff resources, they do not currently implement many of the high-level analytical processes afforded by the HASTUS system that would improve the development of service trips, blocking, and runcutting and would result in more efficient operations. Currently, the time Transit's buses are waiting at the end of routes is higher than at other transit agencies in some cases, the amount of time allotted exceeds what is needed for operations. Finally, Transit has a limited working knowledge of HASTUS and has not maintained the system appropriately, consequently, it uses manual processes to build its schedules and assign resources rather than taking advantage of HASTUS' automated optimization functions. There are significant opportunities for Transit to more effectively use their software to improve productivity and service efficiency.

Summary of Recommendations

Transit should develop a plan to implement the schedule efficiency tools related to service development in recommendations B1 a-j. The plan should identify efficiency targets and propose a timeline for putting each tool into operation.

Implementing Service Efficiency Recommendations Would Save Costs and Alter Scheduling Practices

A two-year timeline would be an aggressive, but achievable, implementation target with full savings being realized about a year following completion of this first phase of implementation. It should be recognized that these recommendations, if adopted, would not be a one-time change, but would alter Transit's scheduling practices. The recommended practices would continue to be employed as part of all future service changes.

Chapter 2

- Expand and implement operating cost efficiency metrics
- Publish and monitor scheduling standards

Chapter 3

- Periodically apply global optimization analysis including the deadhead matrix.

Chapter 4

- Regularly conduct systematic round trip cycle time analysis

Chapter 5

- Implement advanced vehicle utilization (blocking) techniques
- Apply advanced operator assignment (runcutting) techniques

Chapter 6

- Ensure accurate HASTUS calibration
- Update and maintain operating cost assumptions
- Ensure accurate HASTUS data field use
- Ensure that schedulers have the knowledge to fully utilize HASTUS

Over Time, Implementing All Recommended Scheduling Tools Could Save up to $23 Million Per Year in Operating Costs

Each of these tools is discussed in the following report in greater detail. They build upon one another to result in significant ongoing/annual efficiency and cost savings for Transit. The timeframe for realizing these savings will depend on the aggressiveness of Transit's implementation timeline. The estimated annual savings following full application of these scheduling tools to Transit's service development process ranges from $15.7 million to $22.9 million, depending upon the amount of savings that are generated from the round trip cycle time analysis, discussed in Chapter 4. In the following exhibit you will not see each recommended efficiency tool listed separately because some tools are building blocks of others and savings from each individual tool are not cumulative. The dollar amounts shown below summarize the savings that could be realized if Transit implemented all of the recommended scheduling efficiency tools. Also, as noted above, these savings will take some time to realize, and should not be considered as an amount that can be cut from Transit's 2010 budget without impacting service. Savings will be realized incrementally with the full savings one year after full implementation of the recommendations.

EXHIBIT A

Summary of Savings From Implementation of Recommendations

Tools to Achieve Schedule Efficiency	Possible Annual Savings
Round Trip Cycle Time Analysis (Ch.4)	$12 million - $19 million
Advanced Blocking Techniques (Ch.5)	$0.7 million
Advanced Runcutting Techniques (Ch.5)	$3 million
Total Annual Savings	**$16 - $23 million**

SOURCE: King County Auditor's Office

[Blank Page]

2 STRATEGIC APPROACH TO SCHEDULING

Chapter Summary

This chapter discusses the opportunity to integrate strategic scheduling practices that improve the efficiency of bus schedules and operations. We recommend that Transit establish, monitor, and communicate efficiency indicators and develop and apply service development standards.

Performance Metrics

Transit Tracks Some Metrics but Should Fully Integrate Metrics that Measure Progress Towards Achieving Scheduling Efficiency Standards

Although Transit tracks metrics related to system reliability and efficiency, it does not fully integrate efficiency targets into its planning processes. Performance metrics are not are not fully used by schedulers to monitor progress toward ensuring that resources such as vehicles and operators are allocated in ways that achieve lowest possible costs. Performance metrics are commonly expressed in terms of ratios and are trended from one schedule production to the next. Efficiency ratios, when utilized over time, ensure that incremental schedule changes do not degrade the cost-effective allocation of resources across the system.

Transit currently tracks key performance metrics such as on-time performance and adherence to budget. Scheduling staff focus on metrics related to operational reliability and available service hours, which are not efficiency metrics. These are appropriate and constitute valuable management tools; however, because efficiency metrics are not in place and tracked, incremental scheduling changes are not evaluated according to their impact on the efficient allocation of vehicle and operator resources.

One observable result of not using and tracking efficiency metrics during the scheduling process is that Transit has maintained high levels of recovery time, the time buses wait at the end of a route, from one year to another without a mechanism for identifying and rectifying the situation. Transit builds its service budget off a base budget that carries current scheduling practices from one schedule to the next.

Appendix 1 recommends a set of metrics. While each indicator applies to different parts of the scheduling process, some apply to the system as a whole while others can also be applied to individual routes. Together they provide a useful tool for analyzing the efficiency of proposed changes.

RECOMMENDATION B1a

Transit should expand its set of efficiency indicators and goals using Appendix 1 and use them as targets when developing schedules. These goals should be used by management to monitor the performance of the service development group and regularly communicated to decision-makers.

Standards/Guidelines

Although there is an array of documents available to schedulers and service planners that discuss Transit priorities and service best practices, Transit does not have specific documented guidance for service development decisions. Three benefits of formal standards/guidelines are that:

Specific Guidance for Service Design and Scheduling Decisions Would Clarify Appropriate Trade-offs

- Standards/guidelines would establish a framework for making decisions about the trade-offs between efficient operations and other scheduling objectives. For example, service planners might need to make a choice between operating buses 30 minutes apart on a routes versus 35 minutes. The 30 minute gap is more logical for customers, but the 35 minute spacing might be more efficient for operations by requiring one fewer bus to provide the same level of service. Standards/guidelines would assist service planning staff by providing guidance about how Transit expects this decision to be made in varying circumstances.

- Standards/guidelines provide direction about how to utilize ridership and run time data during planning. The types of analysis that are routinely performed, their frequency, and conditions that are analyzed would be identified.
- Standards/guidelines provide accountability and transparency to the stakeholders who fund Transit's services, and along with the performance metrics described above, serve as a basis for understanding the specific costs and rationale for decisions and assessing how efficiently and effectively those funds are used to deliver transit services.

Transit reports that an internal draft of the 10-year strategic plan update includes a work program commitment to compile existing guidance for transit service to be used internally, by the public and by Transit's partners. The draft work program also calls for the development of new standards and guidelines to replace outdated or missing information.

RECOMMENDATION B1b

Transit's planned standards/guidelines document should include be completed, formally adopted, and published, providing a policy guide for Transit staff and reference document for external stakeholders.

[Blank Page]

3 GLOBAL OPTIMIZATION

Chapter Summary

This chapter discusses the opportunity to implement global optimization techniques that consider the system as a whole to improve the efficiency of bus schedules and operations.

Background

Systematic Global Optimization Could Reduce the Number of Buses During Peak Periods, Saving Capital Costs

Transit does not have a systematic global optimization process in place. Global optimization is the process of evaluating the entire transit system to assign work to bases to allow the most efficient utilization of resources. As is common industry practice, schedules at each Transit base are currently developed independently by an assigned scheduler. Typically, schedulers jointly review base assignments during each service change, which occurs three times a year. Their goal is to address the distribution of the fleet based on possible service economies and base capacities. While useful, this process is largely manual and does not provide a comprehensive review of the system that pairs routes together in ways that reduce excessive recovery time. Global optimization that fully utilizes HASTUS and its related modules considers every route in the system at the same time, allowing the scheduler to ask a number of questions and to find options that one might not discover through a mostly manual process. Examples include:

- Which base provides the most cost-effective starting point to and from the ends of the route?

- What opportunities exist to interline[2] routes, thereby reducing costs? Are those routes being dispatched from the same base?
- Can buses that operate only during peak times be parked at another base during the midday to reduce operating costs?
- What is the parking capacity of the base from which this route is dispatched? Are there other routes that should take precedence for being dispatched from this base if that capacity is exceeded?

$459,000 Annually Could Be Saved by Modifying How Routes Are Assigned to Each Base

Global Optimization is a best practice for the efficient scheduling of bus service, offering scheduling efficiencies that may not be apparent when scheduling work is limited to the specifics of each individual base, but becomes visible when the entire system is considered utilizing the full capabilities of the software.

Because global optimization involves systematic evaluation of the entire system, no test at individual bases will necessarily identify all of the economies that may exist for the system; however, a test of global optimization strategies was completed for North, South, and Ryerson bases, looking for opportunities to dispatch buses from one base to provide service on routes that are normally fed from another base. This test identified an additional 12 hours a day and five additional peak buses that could be removed without changing service levels. Expanding the concept to the entire system could achieve savings of $459,000[3] per year. Additional training for operators would be required to ensure that backup drivers were qualified to drive these additional routes.

[2] Interlining is the practice of using the same vehicle or driver on more than one route without going back to the base between route changes.
[3] The $459,000 in annual savings is rolled into the total savings on pg. 7 and should not be added to it. It is included here to show the impact of implementing global optimization.

Deadhead Matrix

Transit's Listing of Travel Times Between Routes in Their Software Is Incomplete, Hindering Analysis of Optimal Routing

Because global optimization systematically analyzes the efficiency of alternative interline arrangements, the deadhead matrix, a listing of travel times between the system's different terminus locations, needs to be completed. As a strategy for reducing overall operating costs, Transit currently interlines many routes. However, the matrix of possible interlines in the HASTUS system is incomplete, precluding a comprehensive automated review of possible interlines. As a test, our consultants developed a deadhead matrix for three bases, looking for opportunities to interline routes as a strategy for achieving economies. This analysis achieved savings of about seven hours per day at the three bases. Expanded system wide, use of an expanded deadhead matrix could achieve annual savings of $275,000.[4] Appendix 2 provides a discussion of tools that may help to accomplish this in a cost-effective manner. In the past, it has been more difficult for interlined routes to stay on schedule than routes that are not interlined.

RECOMMENDATION B1c

Transit should develop a process and procedures for periodic global optimization of its bus system schedule. This should include reviewing and completing the deadhead matrix.

[4] The $275,000 in annual savings is rolled into the total savings on pg. 7 and should not be added to it. It is included here to show the impact of implementing the deadhead matrix during the global optimization process.

[Blank Page]

4 SERVICE TRIP DEFINITION

Chapter Summary

This chapter discusses the process of service trip definition and recommends implementing round trip cycle time analysis.

Background

Definition of service trips is the first activity in the process of schedule production and consists of identifying the time it takes to run buses on routes and determining how often buses will run on each route. Transit systems commonly divide bus service time, or "hours," into three major categories:

Round Trip Cycle Time Analysis Could Result in Savings up to $19 Million per Year

- Run time, when a bus is moving and available to carry riders
- Recovery time, when a bus is waiting at the end of a route[5]
- Deadhead and pull-out/pull-in time, when the bus is traveling from one route to another or to and from the base and is not available to carry riders

Of positive note is the fact that in every route surveyed for this audit, Transit assigned routes to the base where the lowest deadhead costs could be achieved. In all cases, the times scheduled for the deadhead trips were consistent with the amount of time it actually takes buses to complete the trip, and the routing of deadheads appeared appropriate. In addition, Transit makes use of midday storage for some buses at Central Base until they are needed for the evening rush. This reduces unnecessary deadheading.

Schedulers also demonstrated solid methods in building schedules that operate well in a variety of conditions on the street. Each scheduler interviewed displayed a thoughtful, logical

[5] Recovery time is also referred to as layover or break time.

approach to solving schedule related problems within their allotted budget.

Run time

Transit Buses Experience Significant Variability in Travel Times for the Same Trip

Transit buses experience significant variability in run times (the length of time it takes to travel from one end of the route to the other) for the same trip on different days of the week, which will ultimately limit the system's ability to operate consistently. Actual run times were calculated on a sample of fifteen vehicle assignment blocks, identifying the average run time and average variability in run times for each trip. On average, the amount of time provided in schedules was within two percent of average run time, but run times on the same trip varied significantly from day to day. Of 94,245 observations of actual run times, 10 percent were more than one minute early and 13 percent were more than five minutes late (the standard by which Transit and many other transit agencies measure on-time performance). Much of this variation takes place within individual trips where travel times vary from day to day. On sampled trips, there was an observed variance of 43 percent between the maximum observed run time and the minimum run time for the sample, although the vast majority of run times fell within a much smaller range.

Much of the variability in run times likely falls outside Transit's control and could be caused by traffic variation, deviation in boarding patterns, short-term blockages (accidents/roadwork), on-board incidents, or differing driving styles. Transit states that a high number of bus stops on some routes both slow operations and increase schedule variability.

Recovery time

While transit systems work to maximize the proportion of time that the bus is available to carry riders, there are four primary reasons a system will have recovery time built into schedules:

- Provide a cushion to allow the bus to depart on time for the next trip;
- Maintain evenly spaced time periods between buses, often called headway and when set according to an even and easily divisible time, like every 30 minutes, is called a "clock-face headway;"
- Provide time for scheduled transfers between routes for customer convenience; and
- Allow time for operator breaks, which are also required by a collective bargaining agreement.

Transit Buses Wait for 17 Minutes of Every Hour That a Bus and Operator Are Deployed

The ratio of recovery time to in-service time is a common metric for analyzing scheduling efficiency. Six other West Coast transit agencies have a combined average of 21.3 percent recovery to in-service ratio while Transit's[6] percentage of recovery time is 29.2 percent. This is 5.1 to 11.7 percent higher than at the other transit agencies. Exhibit B below summarizes the recovery-to-service ratios at several transit agencies.

EXHIBIT B

Recovery-to-Service Ratio – Transit and Other Systems

Transit Agency	Ratio of Recovery Hours to Service Hours
Phoenix	20.4%
Denver – RTD	24.1%
Portland - Tri-Met	21.3%
San Jose – Valley Transit Authority	23.0%
San Diego – MTS	21.4%
Santa Monica – Big Blue Bus	17.5%
King County Transit	29.2%

SOURCE: Nelson Nygaard

[6] This analysis excludes Transit's Atlantic Base, from which the trolley bus fleet is dispatched. The unique operating challenges of the trolley fleet are addressed in a separate report.

Round Trip Cycle Time Analysis

Transit Round Trip Cycle Times Are Frequently Inefficient

Transit's round trip cycle times are frequently inefficient, meaning that more time and financial resources are used than are required to maintain schedule reliability. Round trip cycle time is the amount of time it takes for a bus to complete one full route cycle (run time plus recovery time). If Transit's round trip cycle times are optimized to the 90 percent level,[7] excess recovery time could be reduced, thus cutting the number of buses needed to meet route requirements and reducing operating costs. A test of round trip cycle time optimization based on 20 sample routes suggests that a range of $12 to $19 million could be saved annually across the system once new analysis processes have been fully implemented.

The implementation and use of round trip cycle time analysis comprises a core recommendation of this audit. The cycle time analysis technique has become possible with the introduction of automated vehicle location (AVL) technologies, which allow transit agencies to compile large databases that accurately portray actual daily run times for every trip in the system during every day they operate. This allows schedulers to conduct statistically valid analyses to determine not only average run times but also the amount of variation in run time that the system experiences and establish the full array of travel time occurrences along a route or route segment.

If actual run times are consistent, the scheduler can be assured that if a bus departs one end of the route on schedule it will likely reach the other end on time. The scheduler does not need to make provision for late trips. On the other hand, if run times are unreliable, the schedule will always be a less consistent predictor of actual performance and the scheduler should properly allow

[7] See detailed discussion of cycle time analysis process on the next page.

additional time – either in the run time or recovery – to compensate for this unreliability.

Recovery Time at the End of a Route Balances Efficiency, Customer Service and Employee Needs

Round trip cycle time analysis is a methodology for determining the amount of schedule unreliability that should be accommodated within a schedule. On any route, some portion of trips will inevitably arrive at the end of the route late. By providing adequate recovery times most transit systems attempt to ensure operators will not arrive at the end of the route after they are scheduled to depart on the next trip.

The first step in a round trip cycle time analysis is identification of an appropriate standard. Several systems identified in this study employ a standard that says either 90 percent or 95 percent of all trips should have sufficient cycle time on their next trip, but this is a local policy decision that should appropriately incorporate local values. As a first step, a system identifies an acceptable level of late operation – 90 or 95 percent in other transit operations. This means that one trip in ten, or one trip in twenty, will arrive at the end of the route after it is already scheduled to depart for the next trip. Given the unreliability of street operations (traffic, accidents, construction, weather, etc.) there will always be a small proportion of trips that arrive after the scheduled departure time for the next trip. The relevant policy issue is determination of what level of such operations that result in late departures on the next trip is acceptable. The more the system opts for reliable operations, the higher the costs of operation.

By using efficient round trip cycle times, the gap between the actual cycle time and minimum cycle time is minimized so the bus can start a new trip instead of waiting for the next scheduled departure. For example, assume that a route operates four trips per hour with a gap of 15 minutes between buses (headway) with departures at /:00 /:15 /:30 and /:45 past every hour and a *minimum* cycle time of 50 minutes. The *actual* cycle time for this route will be 60 minutes because 60 minutes is the closest multiple of 15 that meets or exceeds the *minimum* cycle time of 50 minutes. This route would be said to have excess recovery time of 10 minutes per cycle, per vehicle.

Transit Will Need to Identify in Policy the Percentage of Late Buses They Will Tolerate

The recommended cycle time analysis process identifies the maximum amount of time that is needed to allow for 90 to 95 percent[8] of buses to complete the trip on time. In order to allow consistent service frequencies, for example every 30 minutes, actual recovery times may be longer than is dictated by the cycle time analysis. In this sense, cycle time analyses identify a minimum recovery time, which may be lengthened by other factors.

A test cycle time analysis was completed for a sample of 20 Transit routes to calculate the minimum recovery time and minimum number of buses needed to allow 90 percent of buses to complete their roundtrip plus recovery in the scheduled amount of time. Exhibit C below shows an opportunity to reduce the number of buses needed for the sample routes, which translates to $2 million annual savings in operating costs.[9]

EXHIBIT C
20 Sample Routes – Number of Buses Required

	AM Peak	Midday	PM Peak
Current Vehicles Required	111	94	121
Vehicles Required Using Cycle Time Analysis	106	85	116
Bus Savings Using Cycle Time Analysis	**5**	**9**	**5**

SOURCE: Nelson Nygaard, Courval Scheduling

[8] There is currently no industry standard for an appropriate percentile to use as a minimum, each agency targets performance statistics based on operational goals and budgetary constraints and the determination of an appropriate target reliability percentile is a policy decision to guide schedulers. Individual agencies often use different minimum percentiles for different types of routes.
[9] The $2 million in annual savings is rolled into the total savings on pg. 4 and should not be added to it. It is included here to show the impact of implementing cycle time analysis.

Although the actual systemwide savings resulting from cycle time analysis cannot be calculated precisely since it is extrapolated from a sample, the range of savings is from a low estimate of $12 million per year to a high estimate of nearly $19 million per year with the most likely achievable annual savings around $16.5 million.[10]

Allowing Efficient Bus Spacing That Is Less Intuitive to Customers Could Result in Savings Beyond the Estimated $12 - $19 Million per Year

Additional savings would be possible on the routes sampled if the routes scheduled using regular spacing between buses, e.g., every 15 minutes or every 30 minutes, are redesigned to allow for spacing that is more compatible with actual trip cycle times. This approach could reduce the ease of use for customers in some circumstances, for example, knowing that the bus comes at 1:00, 1:30, 2:00 and 2:30 is easier than remembering that the bus will arrive at 1:00, 1:35, 2:10, and 2:45, but the trade-off is for Transit to "spend" an additional bus and operator to schedule a route every 30 minutes instead of every 35 or 40 minutes. Another option, which requires more analysis and very likely public involvement, is to change the length of the route, most typically shortening it, so that the time required to complete the route is more compatible with an even headway.

[10] Three scenarios, each using slightly different assumptions regarding the level of savings from using round trip cycle times were developed. All three assumed that a 90% standard will be employed.

1. The "High" scenario assumes all routes have the same characteristics as sample routes based on peak bus characteristics. 121 out of 1,010 peak buses were tested in the sample.
2. The "Medium" scenario separately considered peak and midday buses between the sample and the system. 121 peak and 94 midday buses were separately evaluated.
3. The "Low" scenario assumes savings characteristics are related only to routes with all day service characteristics. Only the 94 midday buses included in the sample were evaluated. No allowance was made for expanded services operated during the commute hours.

RECOMMENDATION B1d

Transit should employ a systematic percentile-based cycle time analysis process systemwide. This system should consider both the variation of trip times within a time period (runtime) and time gaps between buses (headways) to determine a minimum round trip cycle time that can be used with confidence for scheduling purposes.

5 BLOCKING AND RUNCUTTING

Chapter Summary

This chapter discusses opportunities to utilize HASTUS to increase the efficiency of the scheduling process during the blocking and runcutting phases.

Background

Transit Currently Uses Primarily Manual Blocking and Runcutting Techniques

Blocking is the assignment of vehicles to a pre-set schedule of trips. A piece of work assigned to one vehicle is called a block. This function can be carried out with the advanced blocking techniques in the HASTUS software, which is generally designed to minimize the number of total vehicle hours from pull-out to pull-in and the number of vehicles needed at the busiest times of day.

Runcutting is the final activity in the schedule production process. It involves breaking vehicle assignments into separate service trips that will be bid on and assigned to individual operators. When the runcut begins, schedules have already been written and combined into vehicle assignments – blocks. During runcutting, blocks are divided into work assignments for operators. Transit utilizes a variety of operator classifications, with varying levels of benefits, who choose individual pieces of work according to the provisions in the labor contract.

Opportunities to Employ Techniques for Blocking Efficiencies

Although Transit seeks incremental manual efficiencies in blocking, the advanced blocking features contained in HASTUS' Minbus module are not being fully employed to create efficient vehicle schedules. Utilizing HASTUS' Minbus module in this process will improve scheduler productivity and increase the amount of the system to be reviewed for efficiency during each schedule change.

$0.75 Million Annual Savings Could Result From Automated Blocking Techniques

By running simulations on three bases using Minbus, total platform hours were reduced by 1.6 percent to 1.8 percent for an average weekday. The savings in peak vehicles that resulted from using Minbus ranged from 30 to 40 buses depending on the scenario considered, an annual savings of $735,095 in operating costs. In addition, by requiring fewer buses at peak times, capital costs of procuring buses would be reduced.

RECOMMENDATION B1e

Transit should utilize HASTUS' Minbus module to implement scheduling procedures that assign vehicles to service trips most efficiently.

<u>Opportunities to Employ Techniques for Runcutting Efficiencies</u>

Transit is not achieving the most efficient run cut, and, like blocking, the runcutting practices currently used by Transit are primarily manual and incremental.

$3 Million per Year Could Result From More Efficient Runcutting

The HASTUS CrewOpt module is a tool for efficiently assigning operators to blocks. CrewOpt employs a complex methodology to minimize total costs and must be accurately programmed in order to produce accurate and efficient results. Transit does not currently utilize this module when runcutting. To test whether Transit would be able to create a more efficient runcut by employing CrewOpt, our consultants developed a runcutting model in order to analyze all bus runs and operator assignments at three bases. Exhibit D shows applying the model at three bases resulted in reductions of about 1.6 percent of total daily operator cost.

EXHIBIT D
CrewOpt Scenario Results Weekday Daily Statistics

Base	Current	Modeled	Weekday Savings	%	Days per Year	Annual Savings
South	$124,881	$122,712	$2,169	1.7%	255	$553,095
North	$79,460	$78,644	$816	1.0%	255	$208,080
Ryerson	$98,856	$96,872	$1,984	2.0%	255	$505,920
Total	$303,197	$298,228	$4,969	1.6%		$1,267,095

Note: Only weekdays were included in this analysis.

SOURCE: Nelson Nygaard, Courval Scheduling

On an annual basis, applying these savings to the sample of three bases modeled results in a savings of $1.3 million per year. With the caveat that this analysis was not conducted on the system as a whole and is highly dependent on careful control of the ratio of full-time to part-time work that must also conform to the labor agreement and actual availability of personnel, if extended to the entire system, the optimization could yield savings of as much as $3 million per year.

RECOMMENDATION B1f

To develop the most efficient runcut, Transit's HASTUS CrewOpt module should be utilized rather than the current manual runcutting process.

[Blank Page]

6 SCHEDULING SOFTWARE

Optimal Use of Scheduling Software Can Achieve Benefits

Chapter Summary

This chapter identifies opportunities to make more optimal use of the scheduling software, HASTUS, and its modules through accurate calibration, updating cost assumptions, using data fields appropriately, and training.

HASTUS Calibration

Calibration is the process of customizing the HASTUS software to local operating conditions and collective bargaining agreement conditions. Calibration is done by the software vendor and/or the transit entity. It entails the development of rules that work as part of the HASTUS scheduling system's automated scheduling functions to:

- improve scheduler productivity by further automating manual scheduling practices;
- improve the efficiency of a schedule by creating more cost-effective work; and
- allow schedulers to conduct scenario testing to quickly evaluate multiple approaches to solve a problem with the fewest resources.

HASTUS' Minbus and CrewOpt tools utilize a combination of hard and soft rules. Hard rules are absolutes, things that must be done or conditions that cannot be violated such as labor contract rules. Soft rules are preferences that can be violated if their cost exceeds a target level. For example, hard rules set the maximum shift length allowed under the contract for each type of operator. In contrast, a soft rule might say that an operator assignment should be limited to nine hours, i.e. incur one hour of overtime, unless doing so would create the need for an additional operator. Both types of rule are needed for HASTUS' tools to provide useful and cost efficient solutions.

Transit has not calibrated HASTUS' Minbus and CrewOpt tools' hard and soft rules to Transit's unique contractual and operating needs and priorities. As a result, schedulers do not have advanced knowledge of using or maintaining HASTUS' Minbus and CrewOpt tools and have limited insight into the tools' use of information within algorithms. All rules guiding the algorithms have been expressed as "hard" rules – even those rules that are unrelated to the labor contract or operating policy. Transit has not set "soft" rule preferences that would typically guide the software to produce an optimum set of work assignments that balance a large set of conflicting goals reflecting historic practices of the agency.

Scheduling Software Has Not Been Adjusted to Fully Reflect Labor Rules

The last CrewOpt calibration was done by the vendor in 2000, which predates the most recent labor contract. As a result, schedulers cannot use HASTUS' CrewOpt module for runcutting because it produces unworkable and unrealistic solutions that are inconsistent with the current labor rules. Since the last calibration in 2000, major changes to operating practices have been handled in-house with consultation from the software vendor, so Transit's new processes are being designed around using manual methods that have limited opportunities for identifying run cut efficiencies. This recalibration of the system will require outside technical resources due to the specialized knowledge of HASTUS and its modules that is required.

RECOMMENDATION B1g

Transit should ensure full calibration of HASTUS to support schedule efficiency and to reduce the time required to produce schedules.

Operating Cost Assumptions

Incorrect Cost Assumptions Have Resulted in Unreliable Software Outputs

Transit's operating cost assumptions specified in HASTUS through rules and parameters were inaccurate at the time of audit work, and Transit does not employ a systematic methodology for identifying the costs that should be programmed into HASTUS. This inhibits the system's ability to achieve the most economical mix of runs. Cost rules, especially operator wage and fringe benefit rates, are not up to date and have not been revised to reflect current conditions. Consequently, HASTUS may produce unreliable results, because rules and assumptions were not current and accurate. Transit scheduling staff do not take full advantage of the HASTUS system's automated runcutting capabilities, citing these unreliable results produced by the system and instead rely on manual run cutting routines in areas where automated capabilities are available.

Regular review of cost information will help to ensure that outputs from the HASTUS software are reliable.

RECOMMENDATION B1h

Transit should develop a systematic process for ensuring that accurate costs are programmed into HASTUS and ensure that it is updated on a regular basis.

HASTUS Data Fields

HASTUS data fields have not been maintained, or in some cases, the data fields have been reallocated for unrelated purposes. The HASTUS software requires maintaining accurate data, located in appropriate fields, to provide accurate and meaningful results.

Some examples of inappropriate data field usage include:

- The "Route Groups" field in HASTUS is typically used for controlling optimization of automated blocking solutions. It has been reallocated as a field to enter data related to ORCA, the new fare payment system. This makes controlling optimization of automated blocking solutions in the Minbus vehicle blocking algorithm ineffective, thus eliminating Transit's ability to define sets of routes that should be considered together or separately during blocking.

Software Data Fields Have Not Been Properly Utilized

- Contractual minimum recovery times are generally used instead of performance-driven minimum recovery times in data fields referenced by the software for automated blocking optimization. This means that the interactive "Create Blocks" command and the automated Minbus optimization tool do not produce realistic and reliable schedules. The minimum recovery is seldom the amount of recovery needed for reliable operations. This forces schedulers to resort to time consuming and less efficient manual methods of blocking schedules.
- The "Vehicle Group" field governs the types of vehicles that may be assigned to a trip on a particular route. Up to three vehicle groups for each trip are allowed by the system. At Transit, vehicle groups have not been applied to all routes or trips to guide the blocking optimization process. Vehicle group preferences are applied neither at the trip level nor at the route level, thus the interactive "Create Blocks" command and the automated Minbus optimization tool produce unacceptable results by mixing routes with different vehicle types within one bus assignment for a given day. This forces schedulers to resort to time consuming and less efficient manual methods of blocking schedules.

Schedulers cannot take full advantage of interactive features or automated optimization features without maintaining accurate data in data fields. This means they cannot use interactive features to improve their own productivity or to improve the efficiency of the schedules they produce, and instead employ manual scheduling practices.

RECOMMENDATION B1i

Transit should maintain accurate data in HASTUS data fields, including restoring algorithm-related data fields to their intended use and creating new user-defined fields as needed for external systems; populating minimum recovery durations for each trip with performance-driven minimum recovery (using the results of cycle time analysis described in Chapter 4); and populating allowed vehicle groups for each trip.

HASTUS Training

Transit's Manual Approaches Provide Limited Opportunities for Maximizing Efficiency and Productivity

Transit schedulers have a limited working knowledge of some modules in HASTUS and currently use manual scheduling approaches in place of automated scheduling solutions, limiting opportunities to make Transit operations optimally efficient and to increase schedulers' productivity. The root cause of many of the scheduling issues facing Transit is a lack of training in use of scheduling software to work faster and to build more efficient schedules. The schedulers cited a lack of resources to implement significant or large-scale changes for a given schedule production. In interviews, schedulers showed a genuine desire to produce a quality product and acknowledged that they lack the requisite training.

Schedulers and service planners should understand the mathematical relationship between minimum cycle times, headways, route length, and the corresponding number of buses required; understand how to enhance efficiency through use of Automatic Vehicle Location (AVL) generated operational data utilizing systematic statistical analysis, are skilled in the use of HASTUS interactive and automated features so they can produce schedules faster and meet service efficiency objectives established by Transit management; and utilize HASTUS to the full extent that its modules allow.

RECOMMENDATION B1j

Transit should ensure that service development staff have the knowledge to fully utilize the HASTUS software system.

APPENDICES

[Blank Page]

APPENDIX 1

SCHEDULING PERFORMANCE METRICS

Items to Be Tracked Without an Established Standard

1. **System Speed** – Total revenue miles divided by total in-service hours. This tracks how fast buses are traveling when they are actually carrying passengers in service.
    a. Track separately for each base and for the system as a whole.
    b. Higher may be better. More importantly, it helps track trends as traffic congestion slows the system down or priority measures speed it up.
2. **Deadhead Miles per Revenue Mile** – Total deadhead miles (base pull-in/pull-out miles plus deadheads between routes) divided by revenue miles. This ratio illustrates whether non-revenue services are increasing as a percentage of the system total.
    a. Track separately for each base.
    b. Lower is generally better.
3. **Peak to Base Ratio** – The maximum number of buses committed to fixed route service during the afternoon commute divided by the minimum number employed during the weekday base period.
    a. Track separately for each base and for the system as a whole.
    b. Generally, this ratio should be about the same as the ratio of system passengers per hour during the peak to the base.

Items to Be Tracked With an Established Standard

1. **Recovery to In-Service Ratio** - This is the ratio of total recovery hours to in-service or "trip" hours for a given schedule. (The in-service trip hours exclude recovery time.) It measures how efficiently trips are hooked together. For example, routes with very short trips of 30 minutes followed by recoveries of 25 minutes will boost this ratio unfavorably. Lower ratios reflect more efficient schedules.
    a. Track separately for each route and each base.
    b. Lower is generally better.
    c. Suggested Standards: Weekday 25 percent, or better, Saturday, 26 percent or better, Sunday 28 percent or better. The goal should be that no base for any day type exceeds 30 percent.
2. **Number (Percent) of Recoveries Exceeding 20, 30, and 40 Minutes** – A physical count of the number of recovery periods that exceed established thresholds. This illustrates the prevalence of recoveries that exceed the standard.
    a. Track separately for each route, base, and for the system
    b. Lower is better.
    c. The target should be that no recovery exceeds 40 (or some identified number) minutes, although exceptions may be warranted when there are few options available and operators are being paid straight through.
3. **Number of Recoveries Less than 5 Minutes** – A physical count of the number of recovery periods that are less than the contractually required minimum. *This should exclude interlines and other situations that are excluded from minimum recovery requirements.*
    a. Track separately for each route and base.
    b. Lower is better.
    c. The target should be that no recovery is less than 5 minutes.
4. **Platform to In-Service Ratio**: This is the ratio of total platform hours – including in-service trips, pull trips, deadhead trips, and recovery time – divided by the in-service "trip" hours. It builds on the first metric (Recovery to In-service) by measuring overall blocking efficiency, including efficiency of how routes are distributed among bases. For example, routes with

very long pull trips from the base will boost the ratio unfavorably. Lower ratios are more efficient.

   a. Track separately for each route and base.
   b. Lower is generally better.

5. **Cycle Time Analysis Check List –** Provides an ongoing accounting of routes that have had cycle time analysis conducted. Once the first round has been completed, each route in the system should be evaluated not less than once every three years. If on-time performance suggests otherwise more frequent is appropriate.
6. **On-Time Performance –** Should be continuously tracked on a route by route, day-type by day-type basis and compared to Transit's adopted on-time performance standard. *Estimated timepoints should be excluded from this analysis.*
7. **Pay to Platform Ratio** – This is the ratio of total pay hours; inclusive of allowances such as sign time, travel time, and premiums such as overtime and guarantee; divided by the total vehicle platform hours. It is used to evaluate the efficiency of a runcut. A ratio of 1.1 would mean that an agency pays an operator one hour and six minutes to operate each platform hour. Lower ratios are more generally efficient, but an important caveat must be considered. Pay hours typically do not include fixed fringe benefit costs. Agencies often boost scheduled overtime hours to reduce headcount and reduce overall payroll costs; this would show as an unfavourable increase in the Pay to Platform Ratio, yet it is an efficient measure, provided the headcount reduction justifies the additional (overtime) pay hours.
   a. Track separately for each base.
   b. Lower is better.
   c. A target should be established that is roughly 1 ½ percent below current practice at each base, e.g., if the current ratio is 1.15, the target would be 1.135. This may need to be adjusted to provide a realistic challenge, consistent with the results of the run cutting simulation.
8. **Part-Time Utilization –** This number should be calculated at each service change. The labor agreement mandates that Transit cannot exceed a 45 percent part time to all operator ratios.
   a. This is a system wide measure.
   b. Under the current cost structure, higher, up to 45 percent, is generally better.

APPENDIX 2

PROPOSED METHODOLOGIES FOR ENHANCING SCHEDULING PRACTICES

While manually blocking a schedule, Transit scheduling staff make many important qualitative trade-offs that involve how trips should hook together to protect reliability of operation and to respect quality of life objectives for the operators who will eventually drive those pieces of work. Their primary consideration is to ensure operators are given sufficient opportunity for recovery while meeting desires of service planners for evenly spaced clock headways.

Identifying what constitutes “sufficient” recovery is currently a subjective component of the manual scheduling process. This should become a more objective component of the process through analysis of historic operating data. In addition, HASTUS has several blocking capabilities that Transit does not fully utilize.

Employ Minbus to Assist in Vehicle Blocking

Transit’s reliance on manual blocking means that the fleet optimization potentials contained within Minbus are not utilized. Each of the following practices that follow will allow Minbus to optimize the fleet assignments.

Employ Dijkstra's Algorithm to Expand the Deadhead Matrix

Dijkstra’s Algorithm, is a graph search algorithm that solves shortest path problems, producing a shortest path tree. This algorithm is only used when values between points have not yet been entered into the deadhead matrix. We are recommending that Transit complete the deadhead matrix. Employing Dijkstra’s Algorithm is an interim measure. For a given set of terminals and known deadhead values, the algorithm finds the shortest path between a terminal and every other terminal based on existing deadheads. For example, if the deadhead from A to B and B to C are known, an absolute upper bound for the deadhead from A-C can be obtained by summing (A-B) and (B-C). If A-C is defined as an existing deadhead in the matrix, the upper bound (A-B) + (B-C) can be used as a cross-check mechanism to ensure coherence between defined deadheads. If A-C is not defined in the deadhead matrix, then the upper bound (A-B) + (B-C) can be used as a conservative default value for the A-C deadhead.

In HASTUS, this algorithm is implemented through the “Calculate from Matrix” command. It has the potential to greatly expand the deadhead matrix, and potentially identifying appropriate deadheads that would otherwise be missed.

Employ Vehicle Groups

HASTUS allows individual trips to be assigned to up to three different vehicle groups. When a trip has one vehicle group defined, it must be blocked using that single mandatory vehicle group. Trips with more than one vehicle group are said to have “alternate” vehicle group options. Note that, when HASTUS links trips together, the most restrictive vehicle group specifications are retained at the vehicle block level. Such fleet optimization gives full consideration to the fleet size of each group, as well as to base capacities.

Make more Extensive Use of Minimum Recovery Time

Minimum recovery durations are not widely used in the current production schedules. Instead, schedulers consider recovery targets based on their knowledge of the route as they manually block trips together according to that knowledge. Zero-value minimum recoveries are often used by schedulers in current production schedules for through-routes downtown that function as if they were one single trip.