DRAFT FOR COMMENTS
- Please do not quote -

BANK OF CANADA
DEPARTMENT OF BANKING OPERATIONS

# Examining the Balance between Risk and Efficiency in Canada's LVTS: A Simulation Approach

Neville Arjani
narjani@bank-banque-canada.ca

5 August 2005

**Abstract**

(... to come)

## 1.0 Introduction

A well-functioning large-value payment system (LVPS) is an integral component of any advanced financial system. In a market economy such as Canada's, virtually all economic (real and financial) transactions ultimately involve a transfer of funds between a buyer and a seller. LVPS provide the electronic infrastructure necessary to facilitate such an exchange of funds between financial institutions in order to discharge large-value obligations on behalf of their own business and that of their customers. Daily payment throughput in LVPS is substantial and many of these payments are time-critical.[1] Key linkages are maintained between these systems and the broader financial system, as LVPS are typically used to settle final funds positions in other important national and international payment clearing and settlement systems, including securities, foreign exchange and retail payment systems. In addition, LVPS are used in the transfer of large-value government transactions, and generally serve as the primary environment for the daily implementation of monetary policy. For these reasons, LVPS are often designated as having the potential to pose systemic risk and therefore fall under the oversight of central banks.[2]

This paper attempts to further understand a fundamental tradeoff in the design and implementation of LVPS – that between settlement delay and intraday liquidity – with specific application to Canada's Large-Value Transfer System (LVTS).[3] Settlement delay refers to a potential time-lag occurring between the intended submission of a payment to the LVPS and when payment finality is achieved, i.e., when funds are exchanged between participant banks on an unconditional and irrevocable basis in order to discharge the payment obligation. Following discussion in BIS (1997), it is argued that settlement delay, as described above, is the primary source of settlement risk in LVPS. Settlement risk is defined here as the risk that final settlement of an individual transaction does not take place as expected. Given the high-value and time-criticality of many payments flowing through LVPS, the cost to participants associated with settlement delay can be large. In addition, the existence of settlement delay may exacerbate potential financial losses associated with other risks in LVPS, such as operational risk.

The design of Real-Time Gross Settlement (RTGS) and RTGS-equivalent systems, such as the LVTS, potentially eliminates settlement delay, as described above. Nonetheless, it is argued that some delay could still exist in these LVPS due to constraints on participants' ability to access intraday credit as a means of funding for outgoing payments. A key source of participants' intraday liquidity, the provision of intraday credit in LVPS is typically subject to collateralization and/or net debit caps. Where sufficient intraday funds are unavailable in a participant's account at the time that it intends to submit a payment to the LVPS for processing, the payment will become

---

[1] In Canada, approximately one-seventh of annual GDP value is transferred through the Large-Value Transfer System on a typical day. This same proportion may equal up to one-fifth in some other countries.

[2] LVPS are generally designated as being systemically important where, if the system were insufficiently protected against risk, disruption within it could trigger or transmit further disruptions amongst participants or systemic disruptions in the financial system more widely (BIS 2003).

[3] The LVTS is owned and operated by the Canadian Payments Association (CPA). For a more thorough description of the LVTS, including an overview of the Bank of Canada's multiple roles within the system, see Dingle (1998), and Arjani and McVanel (2005, forthcoming).

queued – either centrally or in the sender's internal queue – and will be released only when the participant's intraday liquidity position improves. Holding all other factors constant, as constraints on the provision of intraday credit are tightened further and participants' access to intraday funds becomes more scarce, an increasing number of intended payments will not meet the necessary conditions for immediate processing and will become delayed in the queue. It is anticipated that the duration of each payment's queued delay will also rise as constraints on intraday credit are increased.

This paper adopts a general analytical framework proposed by Berger, Hancock and Marquardt (1996) (hereafter known as BHM) as a useful starting point for understanding the tradeoff between settlement delay and intraday liquidity in LVPS. Potential improvements in this tradeoff – defined as achieving a lower level of settlement delay for each amount of intraday liquidity – are also sought. Implementing a complex queue-release algorithm within the LVPS central queue may lead to such an improvement. Primarily used as a gridlock resolution mechanism in the past, gains in computing power have made it possible to apply these increasingly complex algorithms to queued payments on a more frequent basis throughout the day. Introduction of these algorithms is expected to lower intraday funding requirements (and costs) for the release of queued payments, resulting in faster processing times and thus lower settlement delay in the payment system.

Using actual intraday LVTS transaction and credit limits data over a 3-month sample period, the paper employs a simulation approach to empirically assess the nature of the tradeoff between settlement delay and intraday liquidity for the Canadian case. Three measures of settlement delay are utilized – daily unsettled transaction value, a daily system-wide delay indicator, and average intraday queue value. Liquidity levels are controlled by altering participants' net debit caps on intraday credit provision. Participants' payment-sending behaviour is treated as exogenous throughout the analysis. Improvements in this tradeoff are attempted through the hypothetical removal of current informal restrictions on use of the LVTS central queue. The LVTS queue employs a complex queue-release algorithm that seeks to partially offset batches of queued payments on a multilateral basis throughout the day. However, under current system guidelines participants' excessive use of the central queue is not encouraged.[4]

It will be shown that a tradeoff similar to that outlined in BHM exists between settlement delay and intraday liquidity in Canada's LVTS. Furthermore, increased use of the central queue results in a lower level of settlement delay for each amount of intraday liquidity according to all three delay measures. However, some potential implications may also accompany this change in LVTS queuing arrangements. The impact on this tradeoff from anticipated changes in participants' behaviour as a result of such modifications is also addressed. It is anticipated that changes in participants' payment-sending and bilateral credit-granting behaviour will produce competing effects on this tradeoff in the

[4] See LVTS Rule #7. There are several hypothesized reasons for this. Perhaps the foremost reason pertains to the issue of whether queue transparency may cause participants to take on external credit risk by crediting clients' accounts with expected incoming funds, prior to these payments actually being received. This was a major concern of central banks at the time the LVTS was being developed (see RTGS 1997; and discussion in Section 5.2).

LVTS. Further development of the simulation model is necessary before the impact of each of these expected behavioural changes can be quantified.

The remainder of this paper is as follows. Section 2 discusses the nature and scope of settlement delay and describes the tradeoff between intraday liquidity and settlement delay in LVPS. An analytical framework for understanding this tradeoff is provided in Section 3, and potential improvements in this tradeoff are also discussed in that section. Section 3 concludes with an application of the analytical framework to the LVTS environment. Section 4 provides an overview of the simulation methodology, as well as a description of the data. Section 5 presents results from the simulations and related discussion. Section 6 offers concluding remarks and some caveats to the analysis.

## 2.0 The Nature of Settlement Delay in LVPS

### 2.1 Settlement Delay as the Primary Source of Settlement Risk in LVPS

Broadly defined as the risk that settlement of an obligation does not take place as expected, settlement risk represents the predominant risk in LVPS. This risk can be examined on multiple levels, ranging from the individual transaction, to a group of transactions, up to the entire system (BIS 2005). The replacement of traditional deferred net settlement (DNS) arrangements for the clearing and settlement of large-value payments with RTGS and equivalent LVPS in most countries means that settlement risk at the system level is no longer a primary concern. For example, in Canada, settlement risk at the system level is completely eliminated since the LVTS is guaranteed to settle under all circumstances.[5] Hence, this paper focuses exclusively on settlement risk at the individual transaction level, which is defined as the risk that final settlement of an individual transaction does not occur as expected. It follows that settlement risk is comprised of both liquidity risk (i.e., final settlement of a transaction does not take place as expected, but does occur at some point thereafter) and credit risk (i.e., final settlement of a transaction never occurs, perhaps as a consequence of a participant default).[6]

This paper adopts a view of settlement risk similar to that in BIS (1997), and argues that the primary source of settlement risk in LVPS is a potential settlement delay occurring between the intended submission of a payment and when the payment becomes final. Figure 1 provides a graphical characterization of this potential settlement delay within the context of the life-cycle of a large-value payment.

[5] The main impetus behind the replacement of traditional DNS systems with RTGS and equivalent LVPS in most countries has been the elimination of systemic risk. Systemic risk is generally defined as the risk that the failure of one participant in a transfer system to meet its required settlement obligation will cause other participants to be unable to meet their settlement obligations when due (BIS 2003). It is argued that the elimination of systemic risk, as defined above, in RTGS and equivalent LVPS is due in large part to the elimination of settlement risk (at the system-level) in these systems.

[6] In this context, the term 'final settlement' refers to the discharge of a bilateral payment obligation between two LVPS participants through an intraday exchange of funds occurring with finality. Consequently, the notion of settlement risk adopted in this paper does not necessarily encompass the *transfer of the settlement asset*, and therefore applies equally to RTGS-equivalent systems where this transfer occurs on a multilateral net basis at the end of the day.

**Figure 1: The Life-Cycle of a Large-Value Payment**

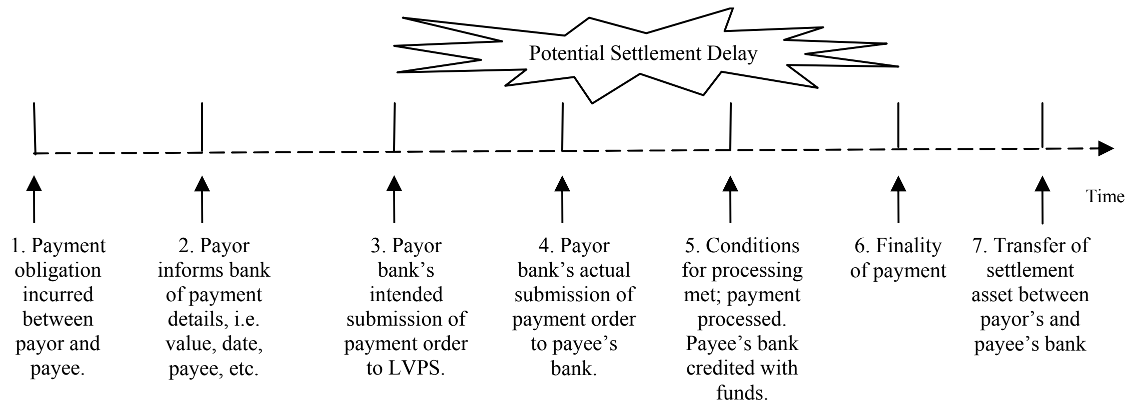


The design and implementation of RTGS and equivalent systems has led to a substantial reduction in (and the potential elimination of) settlement delay, as defined above. Payments are processed in real-time on a gross basis under these system designs and all payments are considered immediately final upon processing. This means the duration between steps 5 and 6 in Figure 1 is reduced to zero under these system designs.[7] Where sufficient intraday funds are maintained by a participant to successfully meet a payment obligation as it becomes due, the duration between steps 3 and 5 is also reduced to zero and settlement delay is completely eliminated. Nonetheless, it is argued that participants in LVPS may confront liquidity constraints throughout the day that disallow the processing of payment messages as expected, resulting in a potential time-lag occurring between Steps 3 and 5. Where intraday funds are insufficient for immediate processing of an intended payment, this payment will become queued – either centrally or held in a participant's internal queue – and will be released only when the sending participant's liquidity position improves to the extent that the payment can be successfully processed.

Some additional points on this interpretation of settlement delay are noteworthy. First, the term 'intended submission' refers to the time that a participant chooses to submit a payment to the LVPS. Thus, the scope of settlement delay adopted by this paper excludes the possibility of participants' strategic payment-sending behaviour influencing the level of settlement delay in the system. Secondly, this interpretation of settlement delay can apply to LVPS both with and without a central queue facility. It is assumed that where a central queue is available, participants will choose to submit all payments to the LVPS at the time they are intended regardless of whether sufficient intraday liquidity is available for these payments to be processed immediately. In these LVPS, the duration between Steps 3 and 4 is assumed to be zero, while a delay may exist between Steps 4 and 5. Conversely, in a LVPS with no central queue, it is anticipated that if an intended payment cannot be processed by the system immediately due to a lack of intraday funds, a sending participant will hold the payment in its internal queue and submit the payment only when its liquidity in the system is sufficient for the payment to be processed

[7] In RTGS systems, the duration between steps 5 and 7 is reduced to zero since transfer of the settlement asset occurs simultaneously with payment processing.

successfully. In this latter case, a delay may exist between Steps 3 and 4, while the duration between Steps 4 and 5 is assumed to be zero.

### 2.2 Possible Consequences of Settlement Delay

Given the time-criticality of certain payments flowing through LVPS, coupled with both the high-speed and high-value of daily payment throughput handled by these systems, the costs associated with settlement delay can be substantial. A participant that is unable to meet its payment obligations when due may suffer the following consequences: reputation damage vis-à-vis its peers in the system; a loss of its clients' business; and even explicit penalty charges imposed by the system operator for delaying time-critical payments. For a receiving participant, not obtaining a payment when anticipated will result in a shortfall in its forecasted intraday liquidity position. Where this participant is planning on using these incoming funds as a source of financing for its own outgoing payments, any delay in receiving a payment means that it may have to incur additional costs in order to replace these funds on short notice. Where it cannot find these funds in time to meet its own obligations, additional settlement delay is created in the system, and the participant is likely to experience similar delay costs as described above. It follows that settlement delay created by one participant in a LVPS could quickly spread to others in the system. Moreover, a comparable disruption to the liquidity position of a receiving bank's client may also occur (where a delayed payment is ultimately intended for this customer), resulting in potentially broader consequences for economic activity.

The existence of settlement delay may also intensify the potential losses associated with other risks in the system, such as operational risk. An operational event will likely have a larger impact the more payments that still remain unsettled in the queue at the time that the incident occurs (Bedford, Millard and Yang 2005; Willison 2004).[8] It follows that the faster payments can be processed with finality within the day, the lower will be the potential financial losses associated with an operational disruption. Also, if faster, more efficient processing of payments helps to encourage greater use of a LVPS relative to less-well-risk-proofed systems, it follows that reductions in settlement delay may translate to lower systemic risk in the broader financial system.

### 2.3 Settlement Delay and Intraday Liquidity in LVPS: A Tradeoff

It is argued that a tradeoff exists between settlement delay and intraday liquidity in LVPS. Intraday liquidity at the individual participant level is generally defined as a bank's ability to meet its outgoing payment obligations as they become due. In RTGS and equivalent LVPS, participants need access to sufficient intraday funds to successfully send payments through the system. Hence, the concept of intraday liquidity could instead refer to a participant's ability to access sufficient intraday funds to meet its outgoing payment obligations in a timely manner.[9] There are two main sources of intraday funds available to participants: 1) funds acquired through incoming payments from other banks

[8] Conversely, an operational disruption could also lead to settlement risk in a LVPS since it may result in a participant's inability to send payments through the system. For this reason, contingency measures are usually available in LVPS for the release of time-critical payments in the event of an outage.

[9] From a purely partial equilibrium perspective, it follows that the earlier in the day that a participant's batch of payments must be processed, the larger the *gross* value of intraday funds it will require early in the day to remain liquid.

due to either normal transaction activity or through an overnight interbank loan arrangement; and 2) funds acquired through an intraday credit extension. The importance of intraday credit as a source of funds for participants deserves emphasis. Martin (2005) argues that the coordination of incoming payments to meet outgoing obligations is often difficult (especially for time-critical payments), and therefore a well-designed LVPS should allow participants to acquire funds when necessary through intraday credit.

Where intraday credit is available to participants on a free (uncollateralized) and unlimited basis, participants can borrow funds any time they need to submit a payment and no time-lag will exist between intended submission and payment processing – settlement delay in the LVPS is eliminated. Although settlement delay ceases to exist in this case, lenders of intraday credit (typically central banks) could face large risk exposures vis-à-vis borrowers, which is not optimal from a public policy perspective. Consequently, intraday credit in RTGS and equivalent systems is not free and unlimited, but rather is often subject to net debit caps, (eligible) collateral requirements which typically entail an opportunity cost, and possibly an explicit interest charge, e.g. the U.S. Fedwire system.

It is argued that constraints on intraday credit provision limit participants' intraday liquidity in a LVPS, increasing the potential for settlement delay in the system. Participants facing these constraints may not be able to acquire sufficient intraday funds to meet their outgoing payment obligations when due, resulting in a delay occurring between the intended submission of these payments and when these payments can be successfully processed. Moreover, as constraints on the provision of intraday credit are tightened, it is argued that the number of payments that are unsuccessful in meeting conditions for immediate processing will rise, coupled with an increase in the duration of queued delay of these payments.[10] In other words, tighter constraints on the provision of intraday credit will result in an overall higher level of settlement delay in the LVPS.

## 3.0 The Delay-Liquidity Tradeoff in LVPS: Analytical Framework

### 3.1 The Delay-Liquidity Efficient Frontier

BHM presents a useful framework for understanding the general tradeoff between risk and cost in LVPS. As the authors argue, this framework can easily be adapted to the study of more specific tradeoffs occurring in LVPS. Following BHM, the tradeoff between settlement delay and intraday liquidity is characterized as a decreasing convex curve in delay-liquidity space (see Figure 2). Each point in the space represents a possible delay-liquidity combination necessary to produce a given level of payment throughput. All points along, and above and to the right of the curve represent feasible delay-liquidity combinations given the current production technology. Points below and to the left of the curve, although preferred, are currently unattainable and can only be achieved through some form of innovation.

BHM introduces the notion of achieving 'technical efficiency' to characterize a LVPS that is operating along the curve. The curve itself is best described as an efficient frontier

[10] For instance, such tightening could take the form of reduced net debit caps, more restricted eligible collateral (which is likely to entail increased opportunity costs for participants), or higher borrowing charges.

in the sense that, holding all other factors constant, settlement delay can only be reduced further for a given level of payment activity by increasing the amount of intraday liquidity in the LVPS. It follows that Point C in Figure 2 represents an inefficient operating choice since both settlement delay and intraday liquidity could be reduced further.

**Figure 2: The LVPS Delay-Liquidity Efficient Frontier**

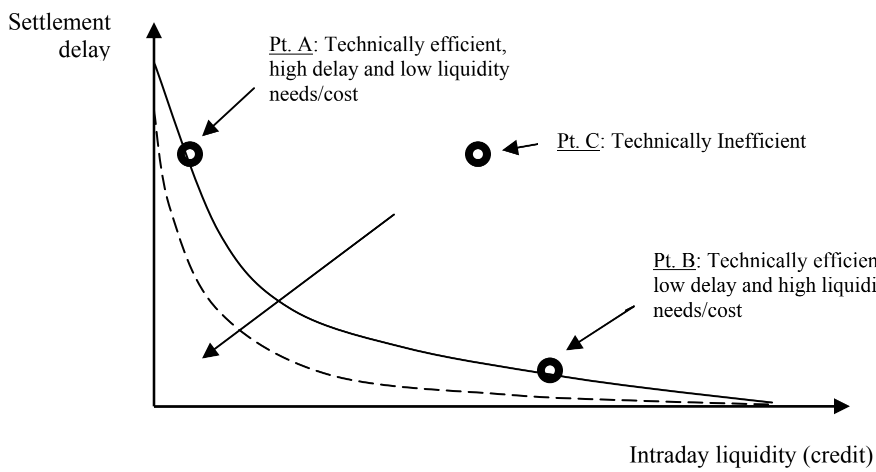


Movements along the frontier from right to left capture the idea that as constraints on intraday credit provision are tightened and participants' intraday liquidity is reduced, settlement delay is expected to rise at an increasing rate. The shape of the frontier is consistent with the assumption of diminishing returns to liquidity. It is anticipated that, at very low levels of liquidity (Point A), a small injection of intraday funding (for example, through reduced constraints on intraday credit provision) will lead to a large reduction in settlement delay since many smaller payments that would otherwise have become queued can now be immediately processed upon submission. As constraints on intraday credit provision are continuously lowered and participants' intraday liquidity is enhanced, more payments will meet the necessary conditions for processing upon submission and delayed finality of these payments will be averted. However, even at high levels of intraday liquidity (Point B), it is expected that a few very large payments will still be delayed where only a substantial injection of intraday funds would allow these payments to be processed immediately upon submission.

It should also be mentioned that efforts to lower the level of settlement delay in a LVPS could increase other risks in the system. For example, lowering constraints on intraday credit provision might result in lenders of intraday funds being exposed to greater credit risk vis-à-vis borrowers. For this reason, it is assumed that there is no 'right' answer as to where along the frontier a LVPS should operate. This will depend on the overall preferences and risk tolerance of LVPS stakeholders.

### 3.2 Introducing an Innovation: Adding a Complex Queue-release Algorithm

Given the potential cost associated with settlement delay in LVPS, it is clear that a downward shift of the delay-liquidity frontier closer towards the origin is desirable to system stakeholders (the dotted line in Figure 2). This shift would mark a definitive improvement in the tradeoff, since it entails a reduction in the level of settlement delay associated with each amount of intraday liquidity for a given level of payment throughput. In this context, a lower level of settlement delay is achieved through faster processing of queued payments or fewer payments entering the queue upon submission, where the latter could occur as a result of the former. Faster processing of queued payments means that intended receivers will obtain expected incoming funds more quickly, reducing the likelihood that their own subsequent outgoing payments will become queued upon submission.

However, points below and to the left of the efficient frontier are not feasible given the current production technology, and a downward shift of the delay-liquidity curve can only be attained through some form of innovation. The question then is, what type of innovation could lead to such improvements in the tradeoff between settlement delay and intraday liquidity in LVPS? One possible innovation outlined in the recent payments literature has been the introduction of a complex queue-release algorithm within the central queue of a LVPS (see McAndrews and Trundle 2001; BIS 2005; Leinonen and Soramaki 1999; Bech and Soramaki 2001; Güntzer, Jungnickel, and Leclerc 1998; and Koponen and Soramaki 1998).

These algorithms are designed to simultaneously search for and offset batches of queued payments, thus serving as an effective coordination device. LVPS participants no longer must wait to obtain sufficient intraday funds for their centrally queued payments to be released from the queue on a gross basis, but rather they need only hold the amount of intraday funds necessary to settle any net debit position resulting from the payment offset. The benefits to LVPS stakeholders from this innovation include lower funding requirements (and related costs) for the release of queued payments, faster processing times for these queued payments, and an overall reduction in average queue length throughout the day. Based on the above discussion, these benefits translate into a lower level of settlement delay in the LVPS, and therefore the introduction of this innovation is expected to result in a downward shift of the delay-liquidity frontier.

The addition of a complex queue-release algorithm does not necessarily represent a new development in LVPS, since these algorithms have been used in LVPS in the past as a gridlock resolution mechanism. However, over the last decade increases in computing power have led to the improved design and more frequent use of these algorithms within LVPS central queues. The complexity of these algorithms has also risen considerably; the choice of full or partial optimization is available and offsetting may take place on a bilateral and/or multilateral basis.

### 3.3 Application: The LVTS Tranche 2 Payment Stream

#### 3.3.1 Description of the LVTS

The remainder of this paper applies the ideas developed in previous sections to Canada's LVTS. The LVTS is an RTGS-equivalent LVPS, where payments are processed on a gross basis in real-time and settlement of the system occurs on a multilateral net basis at the end of the day. The system's risk controls and collateral arrangements, coupled with a residual guarantee provided by the Bank of Canada, provide certainty of settlement for all payments processed throughout the day.[11] This allows funds to be exchanged intraday between participants on an unconditional and irrevocable (i.e., final) basis. Recipients of LVTS payments can make use of these funds immediately upon receipt without any chance of a payment becoming unwound under any circumstances.

The LVTS consists of two payment streams – Tranche 1 (T1) and Tranche 2 (T2) – and participants may use either stream when sending payments through the system. Each stream has its own real-time risk controls. Participants' intraday liquidity in each stream is facilitated by incoming payments received and also by an intraday line of credit. Intraday credit provision in the LVTS is subject to (eligible) collateral requirements and also net debit caps. Eligible collateral consists mainly of government securities, and also high-quality corporate debt. Collateral requirements used to secure intraday credit in the T2 payment stream are more economical relative to those in the T1 stream. Specifically, T1 is characterized as a defaulter-pays payment stream, where intraday credit is secured by eligible collateral pledged by the borrowing participant on a dollar-for-dollar basis. In contrast, T2 is characterized as a survivors-pay payment stream where a collateral pool is used to secure intraday credit such that LVTS participants collateralize the single largest net debit (payable) possible in the system according to well-defined, incentive-compatible risk protocols. It follows that T2 is the dominant stream for LVTS interbank payment activity, representing approximately 87 percent of daily value and 98 percent of daily volume on average since the system began operations in 1999.[12]

The LVTS employs a central queue that stores all submitted payments failing the system's real-time risk controls on a First-In First-Out (FIFO) basis within each tranche type.[13] The LVTS queue is equipped with two payment-release mechanisms. The first mechanism re-submits a participant's queued payments against the risk controls individually on a FIFO basis each time that its intraday liquidity position is increased within the applicable tranche, i.e., the participant receives a payment or its net debit cap is raised. The second release mechanism is an offsetting algorithm that runs at frequent intervals (every 20 minutes) throughout the payment cycle. This complex queue-release algorithm, called the Jumbo algorithm, searches for and offsets full or partial batches of queued payments on a multilateral basis within each tranche type.[14] Payments successfully released from the queue as a result of either of these mechanisms are processed individually by the LVTS as normal.

---

[11] In the extremely remote event of multiple participant defaults in the LVTS, and if collateral value pledged by participants to the Bank of Canada is not sufficient to cover the final net debit positions of all defaulters, the Bank stands ready to exercise its settlement guarantee by realizing on available collateral and absorbing any residual loss.

[12] Most T1 activity is between the Bank of Canada and other participants, and typically involves the transfer of funds related to time-critical payments.

[13] Under current LVTS queuing arrangements, only 'jumbo' payments (i.e. >$100 million) failing the risk control checks become centrally queued in the LVTS. Participants are unable to revoke or reorder centrally queued payments. A payment expiry algorithm is applied to queued payments at regular intervals throughout the day, automatically revoking payments that have been in the queue for greater than 65 min. at the time that it runs.

[14] This is precisely the case for queued T1 payments. For queued T2 payments, the Jumbo algorithm consists of two stages. The algorithm begins by partially offsetting queued T2 payments on a bilateral basis. Once a resulting full or partial batch is determined based on this bilateral offset, full multilateral offsetting is subsequently attempted for this batch of payments in the second stage (see Arjani and McVanel, 2005 forthcoming).

Current LVTS rules state that excessive use of the central queue is not encouraged and the CPA maintains the right to begin charging usage fees if it feels that the central queuing facility is being abused. Participants are able to track their bilateral and multilateral positions in real-time through their internal LVTS workstations and are expected not to submit payments that will fail the risk controls. Anecdotal evidence suggests that participants are adhering to this rule, and only two to five payments are held in the central queue daily. Instead, LVTS participants utilize internal queues to manage submission of their payments to the system. Internal queues are typically equipped with an automated bypass-FIFO release mechanism, which attempts to re-submit a participant's queued payments against the LVTS's risk controls (within its internal workstation) individually on a by-pass FIFO basis each time that its intraday liquidity position is increased.[15] If this process reveals that that an internally queued payment is able to pass the risk controls, it is automatically released to the LVTS for immediate processing.

### 3.3.2 Intraday Liquidity, Collateral and Settlement Delay in the T2 Payment Stream

Since the T2 payment stream is the dominant payment stream in the LVTS, it is the focus in this paper. Intraday liquidity in this tranche is facilitated by T2 payments previously received and also by drawing on a T2 intraday line of credit. This intraday line of credit is subject to both a (indirect) collateral requirement and a net debit cap. Specifically, LVTS participants grant bilateral credit limits (BCLs) to each other, where the value of a BCL represents the maximum bilateral T2 net debit position that a grantee may incur vis-à-vis the grantor at any time during the payment cycle. A participant's T2 intraday credit limit, known as its T2 Net Debit Cap (T2NDC), is calculated as the sum of all BCLs granted to the participant by others in the system multiplied by a system-wide parameter (SWP). The T2NDC represents the maximum multilateral T2 net debit position that a participant can incur during the LVTS payment cycle. The T2NDC of hypothetical bank n (where n = 1,..,N) is calculated as follows:

$$T2NDC^{n} = \sum_{j=1}^{N-1} BCL_{jn} \bullet SWP$$

It follows that two real-time risk controls are applied to payments submitted to the T2 payment stream. A payment will only be processed if it does not result in the sending participant violating either its BCL vis-à-vis the receiver or its T2NDC.

A survivors-pay collateral pool is used in T2 to facilitate LVTS settlement in the event of participant default. Participants are required to pledge T2 collateral equal to the value of the largest BCL that they grant to any other participant, multiplied by the SWP. The value of this T2 collateral obligation is referred to as a participant's Maximum Additional Settlement Obligation, or MaxASO. Essentially, a participant's MaxASO represents its

[15] Under bypass-FIFO, when a participant's intraday liquidity position improves, its first (earliest) queued payment will be re-tried against the risk-controls. If it does not pass, this payment will be by-passed and the participant's second queued payment will be re-tried, and so on. This differs from a standard FIFO set-up, where queued payments are similarly stored in the queue from earliest to latest, however no by-pass function is available when re-submitting payments against the risk controls.

maximum financial loss allocation as a result of another participant's default in the LVTS. Hypothetical bank n's MaxASO is calculated as follows:

$$MaxASO^{n} = \max(BCL_{n,j \neq n}) \bullet SWP$$

The SWP is an exogenous parameter established by the CPA and is currently equal to 0.24.[16] It follows that any change in the SWP will entail a benefit and cost for system participants, holding BCL values constant. In particular, a reduction in the value of the SWP will lower participants' T2 collateral requirement (and related cost) but only at the expense of reduced net debit caps to ensure that risk remains controlled.

Regarding settlement delay in the LVTS, central queuing is seldom used on a daily basis and participants instead hold payments in internal queues until they can be successfully processed by the system. Therefore, settlement delay in the LVTS primarily occurs between a participant's intended and actual submission of a payment to the system. This corresponds with a gap between Steps 3 and 4 in the payment life-cycle depicted in Figure 1.

### 3.3.3 Settlement Delay and Intraday Liquidity in T2: Tradeoff and Improvement

It is hypothesized that a tradeoff exists between intraday liquidity and settlement delay in the LVTS. Moreover, it is believed that this tradeoff shares similar characteristics with the analytical framework presented earlier. In particular, holding BCL values constant and assuming that no migration of T2 payments to the T1 payment stream occurs, reducing the SWP beyond its current value of 0.24 will increase the level of settlement delay in the T2 payment stream. Recall, a reduction in the SWP will directly result in a decline in participants' T2NDC, thus lowering T2 intraday liquidity in the system. Under current queuing arrangements, delayed payments are expected to accumulate in participants' internal queues until sufficient T2 funds are received for these payments to be successfully processed by the LVTS. The magnitude of this settlement delay (in terms of the number of payments becoming delayed and these payments' duration in the queue) is expected to rise at an increasing rate as the SWP is reduced further. Participants will become constrained by their T2NDC more quickly and frequently throughout the day when trying to send payments. In the extreme case, an SWP equal to zero will generally result in a state of payments deadlock where settlement delay reaches a maximum – no participant will have access to T2 intraday credit and therefore will not be able to incur a T2 net debit position. Consequently, no payments will be sent and all will remain unsettled in participants' internal queues until the end of the day.

It has been argued that an improvement in the tradeoff between settlement delay and intraday liquidity can be achieved with the introduction of a complex queue-release algorithm in the central queue. The LVTS already contains a central queue complete with a partial offsetting algorithm. By allowing unrestricted use of the LVTS central queue (and this advanced offsetting algorithm), it is anticipated that participants would no

---

[16] The initial value of the SWP was 0.30 when the LVTS began operations on 4 February 1999. It has since been gradually reduced and has been equal to 0.24 since 16 March 2000. See LVTS Rule #2. The choice of SWP value (i.e., SWP < 1) reflects the effects of multilateral netting (Engert 1993).

longer need to manage an internal payments queue, and instead would submit all payments to the LVTS at the time they are intended regardless of whether these payments could be immediately processed by the system. Essentially this arrangement shifts potential settlement delay in the LVTS from between Steps 3 and 4 to between Steps 4 and 5 in the payment life-cycle outlined in Figure 1. However, potential settlement delay under this alternative central queuing arrangement is expected to be of a lower magnitude compared to under the current internal queuing arrangement, since only under the former regime can a complex queue-release algorithm be applied to queued payments. The key benefit of central queuing compared to internal queuing is that offsetting of queued payments is only possible in the former, and this, in turn, could increase the efficiency of the system.

In the following sections, a simulation approach will be utilized using actual LVTS transaction and credit limits data to shed light on the following questions:

- Under current internal queuing arrangements, does a tradeoff exist between intraday liquidity and settlement delay in the LVTS T2 payment stream? If a tradeoff does exist, is it consistent with the assumptions in the BHM framework?
- Could increased use of the central queue improve this tradeoff? In other words, can the level of settlement delay associated with each amount of intraday credit be reduced for a given level of payment throughput by allowing unrestricted use of the LVTS queue?

## 4.0 Data Description and Simulation Methodology

### 4.1 Description of Data

Three months of LVTS T2 transaction and credit limits data have been extracted over the period July-September 2004. These data represent 64 business days and approximately 1.05 million transactions. This data sample is believed to be representative of normal LVTS activity. Table 1 provides a summary of these transaction data.

**Table 1: Summary of LVTS T2 Transaction Data**

| | Jul 2004 | Aug 2004 | Sep 2004 |
|---|---|---|---|
| **Total Value of T2 Payments ($ billion, CAD)**<br>**(% of LVTS Total)** | 2,283.0<br>(87.8) | 2,203.5<br>(87.9) | 2,446.5<br>(86.3) |
| **Total Volume of T2 Payments**<br>**(% of LVTS Total)** | 349,948<br>(98.0) | 344,357<br>(98.0) | 356,676<br>(98.1) |
| **Daily Average Value ($ billion, CAD)** | 108.7 | 100.2 | 116.5 |
| **Daily Average Volume** | 16,664 | 15,653 | 16,985 |
| **Average Payment Value ($ million, CAD)** | 6.52 | 6.40 | 6.86 |
| **Median Payment Value ($ CAD)** | 42,436 | 40,377 | 45,719 |

An average of $2.31 trillion was sent through the T2 payment stream during each month, in the form of approximately 350,000 transactions. Average daily T2 payments value was highest in September ($116.5 billion) and lowest in August ($100.2 billion). A lower value in August is expected given that the Canadian civic holiday occurs during this month. T2 payments value reached only $6.9 billion on this holiday in 2004.

Payment activity appears to be somewhat concentrated. The Hirschman-Herfindahl Index (HHI) represents one method of measuring the degree of concentration, and is calculated as follows for T2 payment activity over the sample period:

$$HHI_{PaymentActivity} = \sum_{n=1}^{N} \left( \frac{Bank\,n's\,T2\,PaymentActivity}{Total\,T2\,PaymentActivity} \right)^2$$

The value of the HHI will vary between 0.50 and 1/N, where N represents the total number of banks in the sample.[17] Obtaining a HHI value of 0.50 means that payment activity is concentrated only among two banks in the system, while obtaining a value of 1/N (or, 0.0769) means that payment activity is distributed evenly among all banks in the system. The average HHI over the entire sample period is equal to 0.1944 and 0.1813 for T2 payments value and volume, respectively. A HHI value in this range is consistent with payment activity being distributed evenly across a system with approximately 5-6 banks. Indeed, further study of T2 activity indicates that the largest five Canadian banks account for between 85-90% of total payments value and volume, with the remainder of activity divided among the other eight LVTS participants.

### 4.2 Simulation Description and Methodology

The simulation analysis is conducted using the payment and settlement simulator developed by the Bank of Finland (BoF-PSSII). This application is currently being used by over thirty central banks. Version 1.0 of the simulator is employed for this analysis and contains only multilateral credit limits functionality. A revised version of the simulator (Version 2.0) is expected to be released in late-2005 and will include both multilateral and bilateral credit limits functionality. As a result, the methodology in this paper includes the assumption that BCL values remain constant in light of reductions in the SWP and proposed changes to LVTS rules on queue usage.[18] Potential caveats associated with this assumption will be addressed later in the paper. Further, participants' payment-sending behaviour is treated as exogenous throughout the analysis, where the time stamp attached to each payment in the data set serves as the time of intended submission regardless of the proposed parameter changes. Anticipated changes in participants' bilateral credit-granting and payment-sending behaviour will also be discussed later in the paper.

Two separate batches (b = 1,2) of simulations will be run where each batch is intended to replicate a different LVPS design. In particular, batches one and two are designed to replicate the current internal queuing arrangement and the alternative central queuing arrangement in the LVTS T2 payment stream, respectively. Each batch consists of eight individual simulations (s = 1,2,..,8), where each simulation is distinguished by tighter constraints on participants' intraday liquidity. Changes in intraday liquidity are introduced by altering the value of each participant's T2NDC. Since it is assumed that

[17] In this analysis, N=13 because the Bank of Canada does not generally send T2 payments.

[18] This may not be a completely unrealistic assumption. Anecdotal evidence stemming from conversations with LVTS cash managers suggest that credit risk concerns are amongst the major factors when deciding on the value of BCL to grant to a counterparty. Therefore, an increase in a BCL may not be forthcoming despite described changes to system parameters.

BCLs remain constant, a reduction in each participant's T2NDC is achieved by hypothetically lowering the value of the SWP. Specifically, each individual participant n's T2NDC in simulation s is calculated as follows:

$$T2NDC_s^n = SWP_s \bullet \sum_{j=1}^{N-1} BCL_{jn}$$

where $SWP_{1,\ldots,8}$ = 0.24, 0.21, 0.18, 0.15, 0.12, 0.09, 0.06, 0.03.

In specifying the first batch of simulations, the objective is to mimic participants' decision to either submit a payment to the LVTS for processing or hold the payment internally when sufficient intraday funds are unavailable. Settlement delay occurring in this batch represents payments being held internally by participants, i.e., the simulator's queue is used to imitate participants' internal queues. A bypass-FIFO queue-release algorithm is specified to replicate current internal queuing practices of LVTS participants. When this algorithm is applied, a participant's queued payments are re-submitted from the queue and re-tried against the risk controls (in the real LVTS, this occurs within the participant's internal workstation) on an individual bypass-FIFO basis whenever its intraday liquidity position improves. Internally queued payments that can successfully pass the risk controls are assumed to be released from the queue and submitted to the LVTS for processing. In the simulation results for this first batch, settled transactions are assumed to be those that participants were able to submit to the LVTS, while unsettled transactions represent those remaining in participants' internal queues due to lack of intraday liquidity.

Specification of the second batch is intended to replicate a central queuing regime similar to that available in the LVTS. In these simulations, two queue-release algorithms are specified which closely match the LVTS's actual release mechanisms. The first of these algorithms is a FIFO (no by-pass) queue-release algorithm which re-submits a participant's centrally queued payments against the risk controls on an individual FIFO basis each time its intraday liquidity position improves. The second is a complex queue-release algorithm which employs partial offsetting on a multilateral basis and is scheduled to run every twenty minutes, similar to the LVTS's Jumbo algorithm.[19] Settlement delay captured in this second batch of simulations is meant to represent payments being held in the system's central queue, i.e., the simulator's queue is replicating the LVTS central queue. In the simulation results for this batch, all payments in the sample are assumed to have been submitted to the LVTS, and unsettled transactions are those remaining in the central queue which cannot be processed due to a sender's lack of intraday liquidity.

[19] Since bilateral credit limit functionality is currently not incorporated in the simulation application, the partial offsetting algorithm used in the simulations does not exactly replicate the LVTS Jumbo algorithm for T2 payments. Choice of this algorithm when specifying the simulations is meant to capture a combination of features relating to the Jumbo algorithm's two stages for queued T2 payments (see Footnote 14). Despite this limitation, the results generated by the simulations are still expected to be useful and relevant. Further, in specifying this second batch of simulations, it is also assumed that the LVTS's queue expiry algorithm is no longer utilized and *all* payments failing the risk control check become centrally queued (not just 'jumbo' payments).

Three measures of settlement delay are calculated for each simulation within each batch. These measures are intended to capture the daily level of settlement delay associated with each amount of intraday credit provision under both the current and alternative queuing environments described above, holding the level of payment activity constant. They are described as follows:

1. Daily Proportion of Unsettled Transaction Value (PU):

$$PU_t^N = \left( \frac{Value\ of\ Unsettled\ Transactions_t^N}{Value\ of\ Submitted\ Transactions_t^N} \right)$$

An indication of the total value of payments remaining unsettled at the end of each day. This measure is calculated on an aggregate level (i.e., across all participants) for each day t in the sample, where t = (1,…,64).

2. Daily System-Wide Delay Indicator (DI):

$$DI_t^N = \left( \sum_{n=1}^{N} \omega^n \rho^n \right)$$

where $\rho^n = \left( \frac{\sum_{i=1}^{T} Q_i^n}{\sum_{h=1}^{T} \sum_{i=1}^{h} V_i^n} \right)$ and $0 \le \omega^n, \rho^n, DI^N \le 1$

Adopted from Leinonen and Soramaki (1999) and commonly used in payment simulation analyses, this indicator is calculated on an aggregate level and is based on a weighted average of each individual (n) participant's daily delay indicator (ρ). This indicator (and the ratio ρ) can take on any value between 0 and 1, where a value of 0 is achieved when all payments in the day are successfully processed by the LVPS upon intended submission and no settlement delay occurs. A value of 1 is calculated where all payments become queued upon intended submission and remain unsettled until the end of the day. Weights (ω) are based on participants' share of total transaction value over the sample period. Calculation of this measure requires dividing each LVTS business day into T=108 ten-minute intervals (i = 1,..,T). The numerator of ρ represents the sum of a participant's queued payment value (Q) over all T ten-minute intervals throughout the day. The denominator represents the sum of the cumulative value of a participant's submitted payments (V) over all T ten-minute intervals throughout the day. It follows that this indicator is influenced by both the value and delay duration of each payment in the queue calculated for each intraday interval.

3. Average Intraday Queue Value (AQV):

$$AQV_t^N = \left( \frac{\sum_{i=1}^{T} Q_i^N}{T} \right)$$

This is an aggregate measure which calculates the average value of queued payments during each day t. It is found by dividing the sum of total queued payment value (Q) over all T ten-minute intervals on each day by the number of intervals per day (T=108).

**5.0 Simulation Results and Discussion**

**5.1 The Delay-Liquidity Tradeoff in the T2 Payment Stream**

Simulation results for each of the three delay measures are presented in Figures 3 through 5. Two curves are presented in each graph corresponding to each batch of simulations. The curve denoted 'By-pass FIFO' portrays the simulation results estimated under current LVTS (internal) queuing arrangements. The curve denoted 'Partial – No Bypass' depicts results estimated under the alternative LVTS (central) queuing environment.

**Fig. 3: Average Daily Proportion of Unsettled Transaction Value**

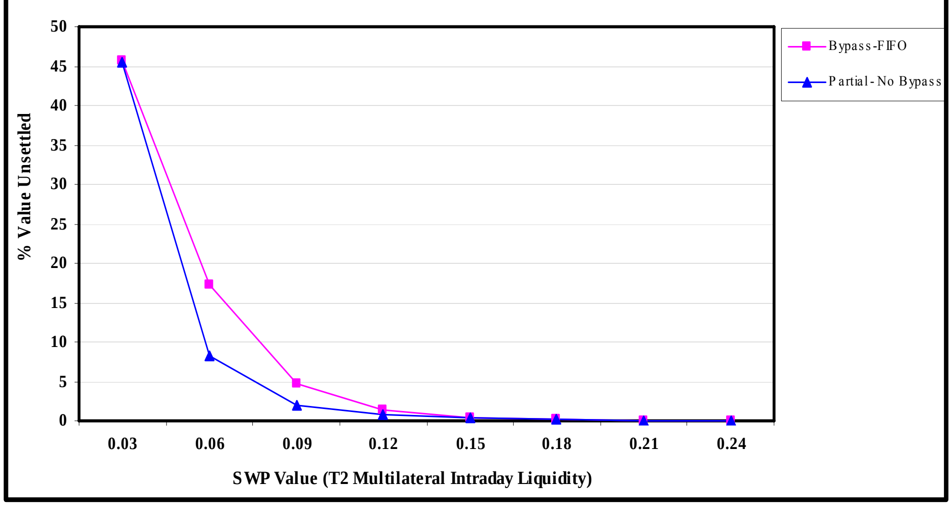

**Fig. 4: Average Daily System-Wide Payments Delay**

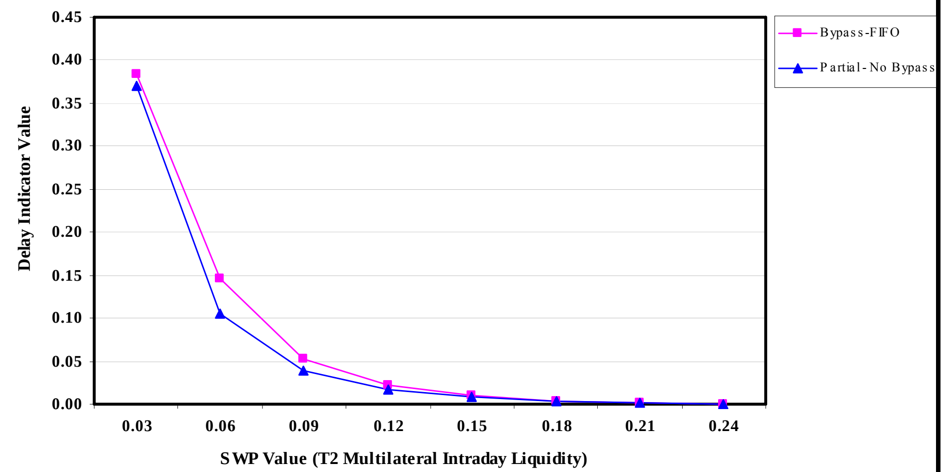


**Fig. 5: Average Intraday Queue Value**

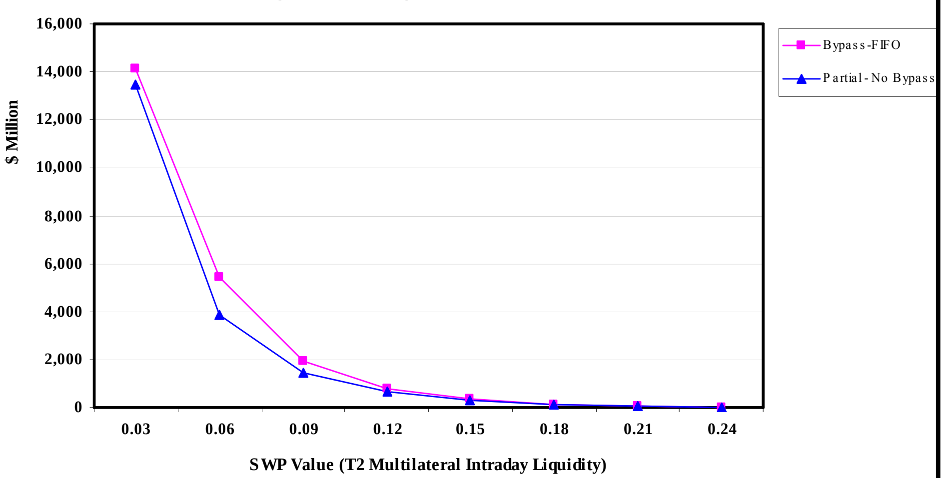


Earlier hypotheses regarding the tradeoff between settlement delay and intraday liquidity are confirmed by the simulation results. Under current LVTS queuing arrangements, a tradeoff exists in the LVTS's T2 payment stream according to all three delay measures. The curve is convex; as intraday credit constraints are further tightened (by lowering the value of the SWP) participants' intraday liquidity becomes more scarce and the level of settlement delay in the system rises at an increasing rate. The slope of this curve increases substantially at low amounts of intraday credit provision.

The introduction of a queuing innovation – allowing unrestricted use of the LVTS central queue – results in an improvement to this tradeoff. According to all three measures, settlement delay associated with each amount of intraday credit provision is reduced following the introduction of a partial offsetting algorithm. The relative benefit of partial offsetting (in terms of reduced delay) increases gradually as intraday liquidity is further constrained. At the SWP value of 0.06, the difference in the level of settlement delay between the two queuing regimes is greatest. When the SWP is equal to 0.06, the alternative central queuing environment is estimated to result in the proportion of unsettled transactions value being reduced by 9 percent or ~$10 billion (Figure 3), average intraday queue value being reduced by 29 percent or ~$1.6 billion (Figure 5) and the system-wide delay indicator being reduced by 28 percent (Figure 4), relative to current queuing arrangements.

Gains from the alternative central queuing design begin to decline when the SWP is reduced beyond 0.06, and the system begins to approach a state of deadlock. When the SWP value is 0.03, settlement delay is only slightly reduced following the introduction of a partial offsetting algorithm, which could mean that participants' intraday liquidity levels are so low that only very small batches of queued payments can be processed each time this algorithm runs. At this level of SWP, close to half of all daily payment value remains unsettled on average under both queuing regimes.

The simulation results can be manipulated to reveal another finding that is closely related to the notion of 'technical efficiency' described in BHM. The above results suggest that settlement delay in T2 increases when the SWP value is lowered from 0.24 to 0.21. However, it remains to be seen whether reductions in the value of the SWP below 0.24 but still greater than 0.21 can be achieved without inducing any further settlement delay in the LVTS. In other words, can a lower amount of T2 intraday credit (and an associated reduction in T2 collateral requirements) be accommodated without increasing the level of settlement delay for payment activity in the three-month sample period, holding all other factors constant? Simulation results suggest that the current value of SWP (= 0.24) is needed to process payments in this sample without increasing the level of settlement delay. A complete discussion of this analysis, including full details of the simulation methodology used, is provided in Appendix I.

### 5.2 Discussion

Some interesting questions emerge from these results which provide direction for future research. First, the benefit of allowing unrestricted use of the central queue is clear from the simulation results – settlement delay is reduced for all amounts of intraday liquidity in the LVTS. Nonetheless, a potential implication of permitting unrestricted use of the central queue pertains to the issue of queue transparency, and specifically whether this reduction in settlement delay could be replaced by an increase in external credit risk taken on by participants. A bank, upon observing an incoming payment in the central queue, may choose to credit its client's account with these expected funds before the payment actually arrives, thus exposing itself to credit risk until the payment is successfully received. LVTS participants have the ability to track expected incoming and outgoing payments in the queue in real-time through their internal participant workstations. Although details regarding client recipients of incoming queued payments are not included in these workstation reports, participants may informally have access to this information.

A second potential implication of increased central queue use relates to LVTS participants' preferences towards reducing settlement delay and lowering T2 collateral requirements. Specifically, it is argued that participants, in granting BCLs to each other, strive to minimize the value of their T2 collateral requirement subject to achieving an established level of throughput efficiency, i.e., an acceptable level of settlement delay. It is likely that payment activity under current internal queuing arrangements may already reflect participants' acceptable levels of settlement delay. Thus, participants may not perceive the benefit of central queuing to be a further reduction in settlement delay, but instead may treat this as an opportunity to realize lower T2 collateral requirements (and costs) while maintaining the same level of settlement delay in the system. This suggests that, under the central queuing arrangement, participants may collectively choose to reduce the BCLs they grant to each other in order to achieve these cost-savings. This reduction in BCLs is expected to continue to the extent that any decline in settlement delay resulting from increased use of the central queue is fully offset.[20]

Following discussion in McAndrews and Trundle (2001), a second potential behavioural response of LVTS participants following this change in queuing arrangements may be to submit more payments to the system earlier in the day, relative to these payments' current intended submission times. The benefits of an offsetting algorithm are expected to increase with the number and value of payments in the queue at the time that it runs. Moreover, anecdotal evidence suggests that participants typically receive information regarding outgoing payment requests well in advance of their intended submission time. Participants' collective submission of as many payments as early as possible to the system under a central queuing regime is anticipated to result in a greater turnover of intraday funds, a lesser need for costly intraday credit, and faster processing of these payments. This reduction in settlement delay is expected to counteract any increase in delay caused by the collective reduction of BCLs as described above. However, as will be discussed in the next section, further research is needed to determine the exact magnitude of these competing effects.

The simulation results also suggest that, holding all other factors constant, the level of settlement delay in T2 will increase only marginally as the SWP is initially reduced from its current value of 0.24. For example, a reduction in the SWP from 0.24 to 0.18 is estimated to increase the average proportion of unsettled daily transaction value by only 0.15 percent under the current queuing regime and 0.14 percent under a central queuing arrangement. Similar results are also observed according to the other two delay measures. Reducing the SWP entails a benefit for LVTS participants in the form of lower T2 collateral requirements (and related cost), as has already been mentioned. Specifically, a reduction in the SWP to 0.18 reduces the aggregate discounted (haircutted) value of T2 collateral required by about $750 million per day on average over the sample period, holding BCL values constant. On one particular day in the sample, the value of T2 collateral required is about $1 billion less when the SWP is equal to 0.18. An interesting area for further research may be to assess whether the benefits of a reduced SWP (in terms of lower T2 collateral requirements) are greater than the associated cost in

[20] Initially, participants are not likely to know exactly how much BCLs must be reduced to achieve the same level of settlement delay under the alternative central queuing regime. Instead, this will be an iterative process that eventually converges to the equilibrium of a perfect offset. In the interim, it may be the case that participants 'overshoot' this target level of BCL reduction, temporarily resulting in a higher level of settlement delay in the system relative to the existing level.

terms of a marginal increase in settlement delay. This entails attempting to quantify the (social) cost of payments delay, and will likely depend on a number of factors including whether these delayed payments are time-critical or not.

## 6.0 Conclusions and Caveats

The objective of this paper has been to gain a better understanding of the tradeoff between settlement delay and intraday liquidity in LVPS, with a specific focus on the Canadian LVTS. Settlement delay represents the primary source of settlement risk in RTGS and equivalent LVPS. Given the high-value and time-criticality of many payments flowing through LVPS, the cost to participants and their customers associated with settlement delay can be substantial. Emphasis is placed on the provision of intraday credit as a source of participants' intraday liquidity in the system. Simulation results indicate that a tradeoff exists in the Canadian LVTS between settlement delay and intraday liquidity, and the nature of this tradeoff is shown to be consistent with the analytical framework proposed by BHM. The results also confirm earlier findings in the payments literature regarding improvements in this tradeoff. Specifically, allowing unrestricted use of the LVTS central queue is expected to reduce the level of settlement delay in the system for all amounts of intraday liquidity relative to current internal queuing arrangements. It was also found that, for the three-month period studied, any reduction in the SWP is expected to increase the level of settlement delay in the LVTS under current queuing arrangements.

Some important issues emerge from these results. Although removing restrictions on use of the central queue is expected to reduce the level of settlement delay in the LVTS, this may also result in a potential increase in the level of external credit risk taken on by system participants. Further, it was shown that hypothesized changes in participants' payment-sending and bilateral credit-granting behaviour in light of the alternative queuing arrangement may fully or partially counteract the initial reduction in settlement delay brought about by increased central queue use. It was also found that under both the current and proposed queuing regimes, an initial reduction in the SWP below its current value results in only a marginal increase in the level of settlement delay in the LVTS, while potentially providing substantial T2 collateral cost-savings for system participants. Further research is necessary in order to quantify whether this benefit is worth the cost associated with increased settlement delay.

The simulation results outlined in this paper are believed to be useful and relevant. However, some caveats apply to these findings. These caveats are discussed here with the intention of motivating further research in the general area of LVPS simulation. The first caveat follows closely with a discussion found in Bedford, Millard and Yang (2005) and relates to the statistical robustness of the findings. The simulation results highlight the estimated impact on settlement delay in the LVTS T2 payment stream as a result of reductions in participants' intraday liquidity over a three-month sample period. Point-estimates of this impact for each amount of intraday liquidity were used to generate the tradeoff curves presented in Figures 3 through 5.

Previous internal research conducted by the Bank of Canada shows that annual LVTS payment activity is affected by specific calendar events and also monthly trends. Consequently, the estimated impact on settlement delay following reductions in intraday liquidity is a random variable that is expected to take on different values based on the dataset used in the analysis. Although using a three-month sample helps to capture the effect of certain monthly and quarterly calendar effects occurring during this period, there is a desire to attain more statistically robust results. For example, it has been observed that the same calendar event may yield a different effect on LVTS payment activity depending on when it occurs throughout the year. Similarly, use of a single three-month sample may not capture the effect that semi-annual and/or annual calendar events may have on the simulation results. Nor will it capture the potential impact of monthly trends in LVTS T2 payment activity.

In order to achieve more statistically robust results, it is suggested that the same simulation methodology be repeated as many times as is feasible using real and/or artificially generated LVTS payment flow data over some fixed sample duration. Grouping the point-estimates of the impact on settlement delay for each amount of intraday liquidity from all of the samples will facilitate generation of an empirical distribution of this potential impact. Figure 6 provides a hypothetical illustration of this result. It follows that the shape of the empirical distribution may be different for each amount of intraday liquidity. For example, the impact on settlement delay may be more volatile and will thus deviate from its mean value more often at lower amounts of intraday credit provision. The shape of the empirical distribution may also change over time.

**Figure 6: Plotting Distribution of Settlement Delay Outcomes**

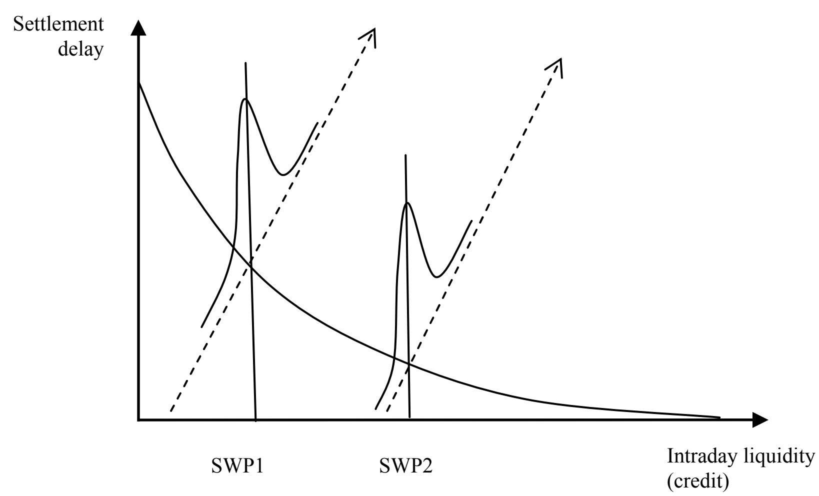


A second caveat applying to this analysis is based on assumptions made in the simulation analysis pertaining to changes in participants' payment-sending and bilateral credit-granting behaviour. Two proposed changes in the LVTS environment are introduced in the paper – a reduction in the level of the SWP and the removal of informal restrictions on use of the central queue. In light of these changes, the current simulation methodology assumes that LVTS participants' payment-sending and bilateral credit-granting behaviour remains unchanged. However, one must question whether this is a realistic assumption. It was argued in the last section that, intuitively, it seems plausible that LVTS participants may react to changes in the LVTS queuing regime in two ways. First, participants may choose to lower the BCLs that they grant to each other in order to reduce their T2 collateral requirements. Secondly, the availability of a central queue is likely to increase the incentive for participants to submit payments to the system earlier compared to under current internal queuing arrangements.

In the context of the analytical framework developed earlier, while the introduction of a central queue with an offsetting algorithm shifts the tradeoff curve downward, holding all other factors constant, the behavioural changes mentioned above are expected to produce further competing effects on this tradeoff in the LVTS. In particular, participants' collective reduction of BCLs is ultimately expected to fully offset this initial reduction in settlement delay for all amounts of intraday liquidity, and can be characterized by a shift in the tradeoff curve back to its original position. At the same time, participants' earlier release of payments to the LVTS relative to these payments' current time of intended submission is expected to result in a further downward shift of the tradeoff curve closer to the origin. However, the ultimate magnitude of this third shift in the curve is currently unknown.

These hypotheses are based solely on intuition, and more sophisticated prediction tools are needed which can be used to develop formal hypotheses regarding the magnitude of potential behavioural changes in response to these alternative system designs. The development of new theoretical and empirical models of the LVPS environment which capture the primary factors underlying participants' payment submission decisions will help in filling this void. Moreover, further expansion of the simulation application is necessary before the effect of these particular behavioural changes can be quantified empirically using actual data, since the current version of the simulator does not include bilateral limits functionality. As mentioned previously, these necessary changes to BoF-PSSII are forthcoming.

A third and final caveat pertains to this same current limitation of the simulation application. Specifically, the absence of bilateral limits functionality creates the possibility that the estimated tradeoff curves provided in Figures 3 through 5 represent a 'lower bound' of the impact on settlement delay resulting from reduced intraday liquidity levels. As T2 multilateral liquidity is reduced (by decreasing the value of the SWP) and payments are delayed, intended receivers of payments won't obtain these funds as expected which may prohibit them from sending their own payments when due. This will certainly result in added volatility in bilateral net positions, possibly to a point where some participants' bilateral net debit positions may actually be greater than the BCLs granted to them. In the LVTS, this cannot occur due to a bilateral risk control test being applied to every payment to ensure that participants are never in violation of their BCL vis-à-vis a receiving participant. Payments failing the bilateral risk control test will be queued until the sending participants' bilateral liquidity position improves. This added delay is not captured in the results generated by the current version of the simulator. Not having BCLs implemented in the simulator forces the assumption that all LVTS payments, when processed by the simulator, have passed not only the multilateral risk control test, but also the bilateral risk control test. It will be interesting to repeat the analysis again with Version 2.0 of BoF-PSSII to compare how much greater is potential settlement delay in the system when bilateral risk controls are also taken into account.

# Bibliography


Arjani, N. and D. McVanel. 2005 (forthcoming). "A Primer on Canada's Large-Value Transfer System." Bank of Canada.

Bank for International Settlements. 2005. "New developments in large-value payment systems." May.

-------------. 2003. "A glossary of terms used in payments and settlement systems." March.

-------------. 1997. "Real-Time Gross Settlement Systems". March.

Bech, M. and K. Soramaki. 2001. "Gridlock Resolution in Interbank Payment Systems." Bank of Finland Discussion Papers. No. 9/2001.

Bedford, P., S. Millard and J. Yang. 2005. "Analysing the impact of operational incidents in large-value payment systems: A simulation approach." In Leinonen, H. (ed.). Liquidity, risks and speed in payment and settlement systems – a simulation approach. Bank of Finland Studies. Pp. 247-274.

Berger, A.N., D. Hancock and J. Marquardt. 1996. "A Framework for Analyzing Efficiency, Risks, Costs, and Innovations in the Payment System." *Journal of Money, Credit and Banking* 28(4): 696-732.

Dingle, J. 1998. "The LVTS – Canada's large-value transfer system." *Bank of Canada Review* (Autumn): 39-55.

Engert, W. 1993. "Certainty of Settlement and Loss Allocation with a Minimum of Collateral." Bank of Canada Working Papers, No. 93-14.

Güntzer, M.M., D. Jungnickel, and M. Leclerc. 1998. "Efficient algorithms for the clearing of interbank payments." *European Journal of Operational Research* 106: 212-219.

Koponen, R. and K. Soramaki. 1998. "Intraday Liquidity Needs in a Modern Interbank Payment System: A Simulation Approach." Bank of Finland Studies, No. E:14/1998.

Leinonen, H. and K. Soramaki. 1999. "Optimizing Liquidity Usage and Settlement Speed in Payment Systems." Bank of Finland Discussion Papers, No. 16/99.

Martin, A. 2005. "Recent Evolution of Large-Value Payment Systems: Balancing Liquidity and Risk." Federal Reserve Bank of Kansas City *Economic Review*. (First Quarter): 33-57.

McAndrews, J.J., and J. Trundle. 2001. "New payment system designs: causes and consequences." Bank of England *Financial Stability Review* (December): 127-136.

Willison, M. 2004. "Real-Time Gross Settlement and hybrid payment systems: a comparison." Bank of England Working Papers, No. 252.

**<u>Appendix I</u>: Is the T2 Payment Stream Technically Efficient?**

The objective of this supplemental analysis is to find the minimum SWP (call this SWP*) necessary to process all payments in the sample without delay, holding all other factors constant. It may be the case that SWP* < 0.24, which means that existing levels of T2 intraday credit, and perhaps more importantly for participants, T2 collateral requirements could be lowered without inducing additional settlement delay for payment throughout in the three-month sample period.

Simulation results produced by BoF-PSSII can provide insight into this issue. Treating participants' payment-sending behaviour as exogenous, a simulation is run using the same sample data but this time specifying unlimited intraday credit. Under this simulation scenario, all payments will pass the risk controls immediately upon submission and therefore no queuing algorithms need to be specified. The daily T2NDC each participant actually needs in order for its payments to be passed without delay can be found from these simulation results, and is equal to the largest multilateral net debit position incurred by each participant during the day. This value is defined as a participant's upper bound (UB) of T2 liquidity. The daily UB of T2 liquidity for each participant can then be used to calculate a value of SWP* that, when multiplied by the sum of the actual BCLs granted to each participant, will produce this UB value. It follows that the highest value of SWP* calculated for any participant on any day is considered the minimum SWP* value necessary to send all payments in the sample through the system without delay. This SWP* can then be compared with the current value of 0.24.

The results from this simulation analysis reveal that on 45 of the 64 days, SWP* reached 0.24 for at least one LVTS participant. This means that the current value of SWP (= 0.24) is necessary for the immediate processing of T2 payment activity during this three-month sample period. Hence, further T2 collateral cost-savings cannot be realized without an increase in the level of settlement delay, holding payment activity constant. The results also indicate that the T2NDC constraint (when SWP=0.24) is binding more often for large LVTS participants (denoted 'B5' in Figure 7). Figure 6 below shows that on 42 days in the sample, at least one of the major Canadian banks reached their T2NDC at some point in the day.

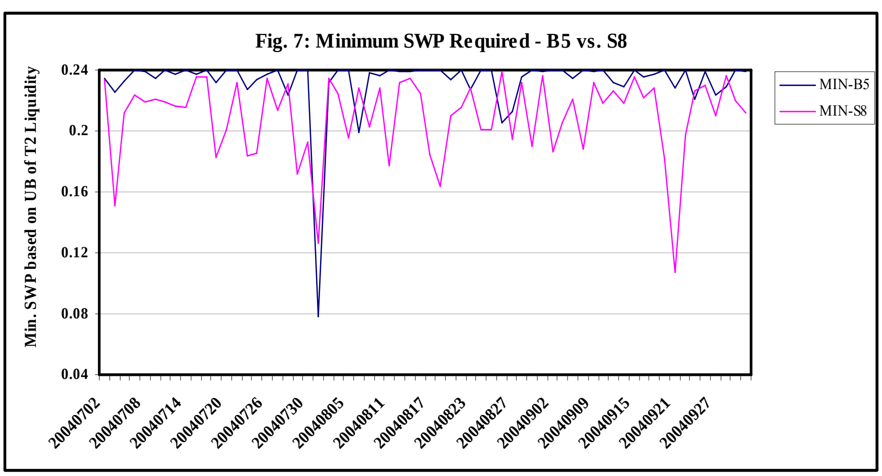

Focusing on the large LVTS participants, the simulation results show that on these 42 days four different institutions bumped up against their T2NDC at least once intraday. One of these participants reached its T2NDC at least once on 37 different days, while the three others reached this limit on 10, 2 and 1 day(s), respectively. It is beneficial to try and understand whether these instances where participants reach their T2NDC are either systematic or random, particularly for the first two participants mentioned above. Specifically, if these participants are reaching their T2NDC at the exact same time of day in each instance, then this may mean that the T2NDC represents a systematic constraint where participants are choosing to structure their payment submission behaviour in a certain way so as to fully exploit their T2NDC value at the same time each day (such as in the early morning hours when the LVTS first opens for general payment exchange). In contrast, if it is found that participants' T2NDC is being reached randomly at different times of the day in each instance, then this may be an indication that this is a random (or true) constraint faced by LVTS participants.

To help in understanding the nature of this constraint (whether systematic or random), the LVTS day is divided into four periods and the time when each participant reaches its T2NDC is located in the simulation results and tabulated. A summary of these findings for the first two large participants discussed above is provided in Table 2.

**Table 2: Percentage of Instances where T2NDC is binding by Time of Day**

| Time of Day | Bank 1 (37 instances) | Bank 2 (10 instances) |
|---|---|---|
| 00:30-06:00 | 0 | 0 |
| 06:00-12:00 | 19 | 0 |
| 12:00-17:00 | 73 | 40 |
| 17:00-18:30 | 8 | 60 |

The results in Table 2 show that these participants are not necessarily meeting their T2NDC at the same time each day, which may be an indication that the current value of SWP represents a random constraint rather than a systematic constraint. Moreover, although not evident based on the results in Table 2, where a high number of instances occur within a certain period (e.g., 27 instances for Bank 1 during the interval between 12:00 and 17:00 hours), these occurrences typically do not take place at the same time within the interval, but rather were scattered throughout the period. This gives further strength to the argument that this may be a random constraint. Of course, more evidence is needed to support this claim. Perhaps further consultations with LVTS cash managers will shed light on this issue.

Some discussion is warranted regarding results for the eight smaller LVTS participants (denoted 'S8' in Figure 7). On only 4 of the 45 days, SWP* reached 0.24 for one of these participants. Further, this occurred for a different participant in each of these four instances. There exist a variety of possible explanations for these results. It may be the case that larger LVTS participants, in sending a higher volume of payments earlier in the day, are 'subsidizing' smaller participants' intraday liquidity in the system, to the extent that smaller participants need to rely less on intraday credit as a source of funding for their outgoing payments. Indeed, SWP* was equal to zero (i.e., no T2 intraday credit was drawn upon) for at least one small participant on 18 of 45 days in the sample. In contrast, this did not occur on any day for large LVTS participants. A second possible explanation could be that, for various reasons, small LVTS participants may tend to bump up against their BCLs far more frequently relative to their T2NDC. Of course, further research is necessary before either of these explanations can be confirmed.