Mit dem Wiederabdruck dieses Beitrags ehren wir das Andenken an Christian Vital, der am 29. August 1997 nach schwerer Krankheit gestorben ist. Sein 20jähriges Wirken für die Schweizerische Nationalbank galt – seit 1986 als Stellvertreter des Vorstehers des III. Departements – insbesondere der Informatik und Fragen des Zahlungsverkehrs. Er war massgeblich an der Entwicklung des Swiss Interbank Clearing Systems beteiligt, das – 1987 in Betrieb genommen – im In- und Ausland hohe Anerkennung als vorbildliches RTGS System gefunden hat. Christian Vital entwickelte sich zu einem weltweit geschätzten und stets hilfsbereiten Experten in Zahlungsverkehrsfragen. Wir danken ihm für alles, was er uns beruflich und menschlich gegeben hat. *Bruno Gehrig**

The Architecture of Real Time Gross Settlement Systems

Christian Vital**

Introduction

The label "Real Time Gross Settlement System" (RTGS) is usually applied to electronic interbank funds transfer mechanisms, which settle large-value payments individually and sequentially via the books of the central bank. Settlement in such systems is final in the sense that it means an unconditional and irrevocable transfer of central bank money from the sending bank to the receiving bank. That is, the receiving bank receives a settlement medium which is free of credit and liquidity risks. The classic example of this arrangement is Fedwire, the wire transfer system operated by the Federal Reserve System in the United States. In net settlement systems, on the other hand, payment orders represent commitments to transfer funds. Settlement of these transfers occurs at discrete-time intervals – commonly at the end of a clearing day – on a net basis. The classic example of this arrangement is CHIPS, the large-value dollar transfer system operated by the New York Clearing House in New York.

The majority of the existing large-value interbank transfer systems are net settlement arrangements. If unprotected, important systemic risks can be created by such schemes. With the spectacular growth of financial markets in the past two decades and the associated growth of interbank payment flows, RTGS mechanisms have attracted – as an alternative – considerable attention from payment system designers in many countries. Prominent examples are the countries of the European Union, the People's Republic of China, Hong Kong, Korea, Thailand and several other countries represented at this seminar. RTGS now seems to have become the norm for the renovation or development of large-value funds transfer mechanisms. The debate over what constitutes an optimal settlement arrangement, however, has only recently been started. So far, it has not led to conclusive results. The distinction between the two polar arrangements is not as clear-cut as one might expect. In reality, the issues are complex. The comparison of benefits and costs is a difficult task, and the results may depend on whether the analysis is made from the point of view of the central bank – the lender of last resort – or from the point of view of individual market participants. Therefore, the question is likely to remain controversial for some time to come.

In what follows, I will discuss key features of RTGS schemes. The focus will be on liquidity and credit risk aspects. The discussion will be based on a concrete example, the Swiss Interbank Clearing (SIC) system, which has been operating in Switzerland since 1987.

* Mitglied des Direktoriums

** This essay was prepared for a conference held at the Hong Kong Monetary Authority on *Global Monetary Systems* on 2-3 November 1995, and was published in the proceedings of that conference in 1996. We are grateful to the Hong Kong Monetary Authority for granting us permission to reprint the essay in this volume. With respect to the original version, the charts have been left out and the figures have been updated to 1996.

A Concrete Example – SIC

SIC is used by banks located in Switzerland for interbank credit transfers of Swiss francs. It is a central facility to which presently 224 participating banks are linked on-line by computer-to-computer connections. The system operates around the clock on bank business days. Settlement is limited to about 22 hours. A value day starts at around 6 p.m. and ends at 4:15 p.m. of the following bank business day. The system was designed for large-value transfers, but it is also being used for small-value transfers. In 1996, 427,000 payments were processed by SIC on an average day and about 1.2 million payments on the peak day.

The daily value of the SIC payment stream was 150 billion Swiss francs on an average day and over 289 billion Swiss francs on peak days. The balances needed for the processing of this payment stream presently amount to around 2.6 billion Swiss francs. This means that a Swiss franc is turned over around 58 times on an average day and over 90 times on peak days.

Demand deposit (or reserve) accounts of the participating banks with the Swiss National Bank (SNB) are administered on the central SIC computer. Debit and credit entries into these accounts are final, that is unconditional and irrevocable. Thus, funds transfers made through SIC represent final transfers of central bank money.

A payment order is settled by SIC if and only if the sending bank has sufficient balances in its account. Overdrafts are not allowed. Uncovered payment orders are automatically held pending in a “waiting queue” until sufficient funds have accumulated from incoming payments and are then automatically released for settlement. The settlement sequence is determined by the priority code which may be attached by the sending bank to the payment order and, for a given priority level, according to the first-in, first-out rule.

Queued payment messages are not released to the receiving bank and may be cancelled any time by the sending bank. The queuing mechanism has no netting or other optimization capabilities. Queue management is only possible through cancellation of a payment order and re-entering it. Payment orders for same-day settlement which are still pending at the end of a SIC day in the waiting queue are automatically deleted by the system and have to be re-entered by the sending bank for settlement on a later value day.

A participant has real-time access to all data available in the system relating to his account. Thus, a participant can monitor settled incoming and outgoing payments, the actual balance of his SIC account as well as queued incoming and outgoing payments.

Since March 1995 the real-time securities clearing system SECOM operated by SEGA, the Swiss Securities Clearing Corporation has been linked to SIC. This link provides a simultaneous delivery-versus-payment (DVP) procedure for securities transfers on a trade-by-trade basis. SECOM earmarks the securities to be transferred and sends the related payment message to SIC. SIC debits and credits the accounts of the buying bank and the selling bank respectively as soon as the buying bank has sufficient funds in its SIC account. The settlement of the transfer is then confirmed by SIC to SECOM, and SECOM transfers the earmarked securities irrevocably. Presently, about 40,000 securities transactions with a value of around 7 billion Swiss francs are settled on an average day through this DVP procedure.

In this outline of the system being operated in Switzerland I have mentioned features which shape the characteristics of RTGS systems. In practice, different designs have been implemented or are being implemented. The variations may result from historical circumstances, differences in the legal system, and so on. In the following I will discuss some variations of features which I believe to be important for the smooth functioning of an RTGS system.

Message Flow

The design of the flow of payment and settlement messages has important implications for the characteristics of an RTGS scheme. Different designs have been proposed in the past few years. The simplest is the so-called *V-structure.* In this structure the payment message is sent by the sending bank to the central bank. After settlement, the central bank sends the payment message to the receiving bank and thereby confirms the settlement of the payment order. Examples of this design are Fedwire and SIC.

In the *Y-structure* the payment message is sent by the sending bank to the node located at the joint of the "Y". Settlement data is then stripped from the payment message and sent to the central bank for settlement. After confirmation of settlement by the central bank, the payment message is delivered to the receiving bank. This design is under consideration, for example, by the Banque de France and by other central banks which rely on the communication services of S.W.I.F.T. for implementing their large-value payment systems.

The so-called *L-structure* is being implemented in the CHAPS system operating in the United Kingdom. In this design, the sending bank first sends a settlement request to the central bank. After confirmation of settlement the payment message is then sent to the receiving bank.

The basic common element in these designs is that the *payment message is released to the receiving bank only after settlement has occurred.* An alternative design is the so-called *T-structure.* In this structure the payment message is sent to the receiving bank already prior to confirmation of settlement by the central bank. As a result, there will be a time lag – the "settlement lag" – between the receipt of the payment message and the confirmation of settlement. Since it may be difficult for receiving banks to distinguish between settled and unsettled payment orders in this scheme, they are likely to put them in the same basket and act upon this information. This would result in similar credit and liquidity risks as observed in unprotected net settlement systems.

Settlement Mechanism

In an RTGS system, settlement of a payment order occurs by transferring reserve balances from the account of the sending bank to the account of the receiving bank. Since, as a rule, central banks do not pay interest on reserve balances, participating banks will economize their holdings of reserve balances as far as possible. The value of the daily payment flow may therefore exceed the stock of reserve balances by a wide margin. For example, in SIC reserve balances are turned over about 58 times on an average day and about 90 times on peak days. This means that incoming payments are an important liquidity source for outgoing payments. If incoming and outgoing payments cannot be synchronized to such an extent that outgoing payments are always covered by existing balances, there are three possible reactions by an RTGS system.

First, the balances needed can be created by allowing the account of the sending bank to be *overdrawn,* that is by extending credit from the central bank to the sending bank. The expectation is that these credits would only exist during short time spans and would be extinguished by balances resulting from incoming payments before the end of the processing day. Whether such daylight overdraft credits should be limited, priced and/or collateralized depends on risk, cost and monetary policy considerations.

Second, the uncovered payment order can be *rejected* by the RTGS system. The sending bank will then have to resubmit it again for settlement at a later time. For large payment volumes this solution presupposes the existence of automated queuing mechanisms within the participant's own systems which interact with the RTGS system until the settlement conditions are fulfilled. The mechanism has the effect of delaying the settlement of an uncovered payment order until sufficient balances have been made available.

Third, an uncovered payment order can be held pending by a *queuing mechanism* implemented on the central RTGS processor until sufficient balances have accumulated.

The common element of the second and third options is the delaying of the settlement of uncovered payment orders by queuing mechanisms until sufficient balances are available for settlement. A combination of one of these options with the overdraft option is, of course, possible and is considered for several schemes which are presently being constructed.

Queuing Mechanism

In what follows, I shall assume a *centralized queuing facility* – the option chosen for most RTGS systems. The processing sequence, the transparency of queued (or pending) payment orders to the sending and receiving banks, and the revocability of queued payment orders are key issues for the design of a queuing mechanism. They have important implications for liquidity and risk management.

(i) Processing Sequence

The liquidity necessary to process a given payment flow crucially depends on the processing sequence.

The simplest choice is the first-in, first-out rule. This rule has the effect that payment orders are settled according to their input sequence. In this scheme, uncovered large-value payment orders may cause small-value payment orders to be queued even though sufficient balances would be available for their settlement. This rule governed the operation of SIC until 1994.

A more flexible approach is to allow a sending bank to assign a priority to a payment order and to process the payment flow according to priorities, and, for a given priority, according to the first-in, first-out rule. This approach gives the sending banks more flexibility in the management of their payment flow. It was implemented in SIC in 1994.

Both methods are static in the sense that, once the sending bank has made its choice regarding the priority and the input sequence, the settlement sequence is determined and cannot be changed. This restriction could be relaxed by a facility which enables a participant to dynamically change the sequence of its queued payment orders, for example, by enabling a participant to point to a payment order which should be given the highest priority. Such an additional option is planned, for example, by The Nederlandsche Bank for its TOP system.

All three methods mentioned so far ensure that the settlement sequence is determined by the sending bank. It is a matter of debate whether changes in this sequence by a third party, that is by manual interventions of the RTGS operator or by an optimization algorithm of the settlement facility, are a desirable further option to increase the efficiency of an RTGS system. One concern is that the central bank could be made liable, if, as a result of reordering actions, payments would fail to be made. Another concern is that intervention by the third party could destroy the scheduling chosen by the participant in order to reach certain goals, for example, in order to reduce the settlement lag in foreign exchange transactions. These concerns could be reduced by restricting interventions to day-end procedures.

(ii) Transparency of Queues

The most controversial issue in the present RTGS discussion is probably the question whether a receiving bank should have *access to queued incoming payment orders.* In SIC, this information is made available on request.

Giving a participant access to queued incoming payment orders may induce him to act upon that information, for example, by honouring a cash withdrawal. If so, he will be exposed to the sending bank in a similar way as in a net settlement system. This is the main argument against making incoming queues transparent to participants. The weight of this argument may depend on the technical implementation of this function. If information on queued payments and the content of the payment message is *automatically released* to the receiving bank, the mechanism would in fact degenerate to a T-shaped RTGS system with the corresponding exposures. If the information is only *available on request,* the situation is different. If the receiving bank acts upon the assumption that a queued payment will be executed, it makes a credit decision which – in a properly managed bank – will be subjected to ordinary credit risk management procedures. The same result would be achieved by using a different communication channel, e.g. the telephone, but in a less efficient way. Indeed, not providing information on pending incoming payment orders would be an incentive for participants to use a separate communication system for this purpose. This, again, would mean that the mechanism degenerates to a T-shaped structure with its risk implications.

A second – in practice more important – aspect to be considered in this context is *liquidity management.* In an RTGS system which does not allow more or less unlimited overdrafts, queued incoming payments are an important source of liquidity. The relevant information is particularly important towards the end of a day if not all queues are empty. The cash manager will have to make his decisions regarding any adjustments of his position by that time at the latest. If the information on the total amount and the originators of queued payment orders are not readily available, errors are likely to occur. The uncertainty would probably also make more or less frequent postponements of final cut-off times necessary in order to give participants the necessary time for collecting the required information through other communication channels (e.g. telephone calls to other participants). As an alternative, the RTGS could be operated in such a way, that, under normal circumstances, the queues would be empty well in advance of the final cut-off time. This goal could be reached by providing sufficient intraday liquidity to the participants, for example, in the form of unpriced daylight overdraft limits.

(iii) Revocability of Queued Payment Orders

The question regarding the revocability of queued payment orders has so far not received a uniform answer. In SIC, a queued payment order can be revoked anytime by the sending bank until the first cut-off without the consent of the receiving bank. Between this first cut-off and the final cut-off the consent of the receiving bank is required. In other schemes queued payment orders can only be revoked by the system operator on request by the sending bank and revocability may be limited to certain conditions, for example, to cases of erroneous payment orders. The third category comprises those schemes which do not allow the revocation of queued messages.

In an RTGS system which queues uncovered payment orders there is the risk that the queued orders will not be executed. In a scheme which restricts or does not allow revoking, receiving banks may rely on the arrival of queued incoming payment orders to a degree which may not reflect the prevailing risks. The revocability rule intends to make the participants of the payment scheme fully aware of this risk.

The revocability of queued payment orders has another important function in RTGS systems which restrict daylight overdrafts. In such systems, *gridlocks* may be a potential problem. This refers to a situation in which payments do not move because they are all awaiting incoming funds. Gridlocks are likely to occur if participants enter amounts which are large in relation to the reserve base available in the system. If revocation is allowed, such payment orders can be revoked and re-entered in two or more portions. This splitting of very large amounts has proved to be an efficient instrument in the hands of SIC participants to resolve "serious" gridlocks.

(iv) Optimization Procedures

Gridlocks are a major concern in RTGS systems – at least if they last until the end of the processing day. The efficiency of the settlement mechanism could be increased by implementing optimization procedures which *reorder* queued messages in order to increase the number or value of settled payments, and it could be increased by *netting* incoming and outgoing payments.

I have already mentioned several concerns regarding intervention in the processing sequence originally determined by the sending bank in the course of the processing day. The question remains whether optimization procedures could help to resolve those gridlocks which remain until or occur near the end of a processing day. The systemic risks associated with such gridlocks can be important. Because of a lack of relevant information and with the exception of very simple cases, the results of an optimization procedure cannot be predicted by a participant. Therefore, rules should be designed in such a way that the participants have incentives to avoid gridlocks. A possible solution is to abstain from optimization procedures and to give participants the option either to take action before the end of the processing day, or to bear the consequences of non-execution of *all* payments involved in the gridlock. This strategy has been adopted for the operation of SIC following the experience that participants started to rely on manual optimization interventions by the SNB in the initial phase of the SIC operation, thereby increasing the risk of gridlocks with important systemic consequences.

Account Structure

With respect to the account structure, two issues have to be considered. The first issue results from the fact that commonly a central bank as well as the banks participating in the payment system have a number of *branches.* If so, should the branches of the bank be permitted to hold accounts with more than one office of the central bank, that is, should a bank be allowed to hold multiple accounts with the central bank for payment purposes? The management of multiple accounts poses cash management problems for the bank and risk management problems for the central bank which are difficult to solve. Therefore, as a rule, centralized accounts are preferred for large-value payment transactions. This does not exclude, of course, that supplementary accounts are maintained for special purposes, e.g. cash withdrawals by a bank branch from the local central bank office.

The second issue is the degree of integration of the accounts used for payment and settlement purposes with the *accounting system* of the central bank. In the Swiss case, these two functions are separated. The accounting system maintains the so-called Master accounts, the SIC system the SIC accounts. Both types of accounts are reserve accounts. At the beginning of a day, balances needed for payment purposes are transferred from the Master account to the SIC account. At the end of the SIC day, the balances of the SIC account are transferred back to the Master account. During the day, balances can be moved anytime from one account to the other. The SIC accounts are used for funds transfers between the SIC participants, the Master accounts for the remaining operations such as cash withdrawals or money market transactions between the SNB and a SIC participant.

As a result of this separation of accounts, the accounting system and the payment mechanism are only loosely coupled. This facilitates the task of designing optimal data processing and communication facilities for the two functions with widely diverging requirements. For example, it very much facilitates the task of providing a 24-hour payment service and the high availability required for a real-time funds transfer mechanism. The disadvantages of this separation are, of course, the same as those already mentioned with respect to multiple accounts. However, their significance is limited since the payment traffic is concentrated on the SIC accounts and only a very small number of transactions are posted to the Master accounts.

Concluding Remarks

There are, of course, many other aspects which are of interest when discussing RTGS systems. Operational issues such as the throughput capacity of the settlement mechanism, the reliability of the system itself as well as the participants systems, which are linked to it and the availability of adequate back-up solutions have important implications for the overall operation of the system. Other examples are the links between the RTGS systems and net settlement systems, the supply of intraday liquidity and related issues. I have concentrated my discussion on those aspects which seem to be controversial in today's RTGS arena and on which we have been able to collect some interesting evidence with the operation of an RTGS system in Switzerland.