New payment system designs: causes and consequences

James McAndrews, Research and Market Analysis Group, Federal Reserve Bank of New York and
John Trundle, Market Infrastructure Division, Bank of England[1]

Payment system design is changing. This article discusses these changes, the pros and cons of different designs and issues raised by the evolution of a variety of so-called 'hybrid' systems. It suggests that there are many common elements in these developments and that outcomes combining lower settlement risk with lower liquidity costs may be possible. But in the end some trade-off may still need to be made between different objectives.

Over the past decade the most popular design for high-value payment systems has changed, reflecting market needs and the concerns of central banks for systemic stability. Before 1990 most systems were designed on a deferred net settlement (DNS) basis but since then almost all major countries have moved to Real-Time Gross Settlement (RTGS) systems for high-value payments. They are now used widely around the world. More recently, however, hybrid systems, which seek to be liquidity-efficient, have been introduced in a small number of developed countries. This reflects attempts to achieve combinations of safety and efficiency that are acceptable both to central banks and to system participants.

This article reviews the main elements of system design and discusses the causes and consequences of the evolution of hybrid systems. It discusses in particular the importance of the status of queued payment messages in hybrid systems and the potential value of queues in offering a signalling mechanism of banks' willingness to pay each other. This can provide incentives for banks to submit payments promptly. The article also draws attention to the different incentives to monitor credit-worthiness in different systems. It suggests that participants' preferred choice of system will depend particularly on:

- the ease and cost of obtaining intraday credit
- the number and size of participants in the system
- the variability of daily payment flows
- the direct costs of participation.

Central banks will continue to focus on whether proposed designs ensure that risks are controlled effectively.

Deferred net settlement systems

In a DNS system, participant banks send each other payment messages during the day and settle the net amounts due at a subsequent specified time, typically at the end of the day. The settlement may occur on a bilateral net basis where each bank pays, or is paid by, each other bank. Alternatively settlement may be on a multilateral net basis where those in overall debit to all other participants pay the net amount due to a settlement agent and the settlement agent pays the banks which, overall, are in net credit. The settlement agent itself needs to be financially and operationally robust as it plays a pivotal role in the system. Often the settlement agent is the central bank. These systems are economical on liquidity use but involve the participant banks implicitly granting each other credit during the day. Should the sending bank fail after transmitting a payment order but before settlement is final, the receiving bank would not receive the funds it expected at the end of day. A failure to settle by one bank in a DNS system could have repercussions on the other banks' ability to settle, potentially compounding the adverse effects of the first settlement failure.

1: The authors are very grateful to Erin Campbell, Market Infrastructure Division, who is the author of the annex and who commented on and contributed significantly to the main article. The views expressed are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Most systems, in G10 countries and elsewhere, were DNS systems until the 1980s. Typically they had few or no risk controls to limit the exposures that could occur during the day. As the values of flows through these systems increased, and as awareness spread of the potential consequences of a default by a participant, central banks called for the tightening of risk controls. The current high levels of turnover in G10 payment systems relative to GDP is shown in Table 1. The Bank of England and the Federal Reserve, among others, drew attention to these risks in the 1980s, and, in 1990, the Lamfalussy Report set out standards for the operation of cross-border interbank netting systems[2].

Table 1: Large value payment systems in relation to GDP (1999)

Country	System	Annual payments turnover/GDP
Belgium	ELLIPS	83
Canada	LVTS	23
France	TBF and PNS	56
Germany	EAF and ELS	31
Italy	BI-REL	28
Japan	BOJ-NET	70
Netherlands	TOP	42
Sweden	K-RIX	51
Switzerland	SIC	111
United Kingdom	CHAPS Sterling and Euro	75
United States	CHIPS and Fedwire	69

Source: BIS and IFS.

Although the Lamfalussy Report was aimed specifically at cross-border systems, in practice it was applied also to domestic systems. It made six recommendations for DNS systems, including a requirement that the system should be able to withstand the failure of the largest net debtor. A protected system could settle in the event of a participant's failure to pay, for example because it held sufficient cash or other suitable collateral provided by the participants. Non-cash collateral would need to be capable of being liquidated quickly so as to obtain funds to complete the settlement. These minimum standards were very influential and have been adopted widely in developed countries. It remains the case, however, that the majority of payment systems around the world are 'unprotected' DNS systems. This is the case in developed countries for many retail systems but is also true of many emerging markets where there is no separate high-value system[3]. As the values going through such systems grow, and as more firms join foreign systems, the weaknesses become an increasing source of concern. Moreover even protected DNS systems complying with the Lamfalussy standards do not avoid all settlement risk. Losses may be shared amongst survivors only up to a limit and the system may still be vulnerable to the failure of more than one participant.

As this discussion indicates, a key feature of DNS systems is that participants must be willing to extend credit to each other. They therefore need to set minimum levels of creditworthiness for their counterparties and to have a means of monitoring each other for changes in credit standing. As a result participants in such systems generally have to meet stringent criteria. Although other institutions may participate on an indirect basis, the limited ability to participate directly in DNS systems may restrict competition between participants.

The risks associated with DNS systems, even those with some risk protection, encouraged the adoption of RTGS, especially for systems dealing with high-value payments.

RTGS systems

An RTGS system eliminates interbank credit risk in the settlement process by providing immediate finality for payments between direct participants. Payments are processed on a transaction by transaction basis and involve a simultaneous debit of the paying bank and credit of the receiving bank in the books of the settlement agent, typically the central bank. If a bank sends a payment order but does not have sufficient funds to complete the transaction, it would not be processed. In some RTGS systems the rejected payment message is returned to the sender for subsequent resubmission. In other RTGS systems it is stored in a central queue until there are sufficient funds on the account for it to be executed. Should there be insufficient funds on the account at the end of the business day the payment would be cancelled[4].

2: See Board of Governors of the Federal Reserve System (1984), Bank of England (1989) and BIS (1990).

3: See the survey in Fry et al (1999) pp. 112–114.

4: A fuller description of RTGS systems can be found in *Real-time Gross Settlement Systems*, BIS (1997).

In the 1990s, all EU and G10 countries (with one exception[5]) that did not already have RTGS systems introduced them. RTGS is now the standard approach for high-value systems used by any country undertaking payment system reform. There are around 70 RTGS systems in operation or planned around the world.

To function effectively, however, an RTGS system requires an adequate supply of intraday liquidity. Each individual bank typically has to make daily payments that are much greater than the amount of liquid funds (such as balances at the central bank) it would typically hold overnight. In most countries the central bank offers to supply banks with the additional intraday liquidity that they need. This credit is usually made available against collateral provided to the central bank (which mitigates the central bank's counterparty risk) and without interest charges. In the United States, however, the Federal Reserve Banks supply intraday liquidity on an uncollateralised basis and impose an interest charge depending on the duration and size of the intraday credit[6].

Liquidity costs

It is clear that explicit interest charges make intraday central bank credit costly and that banks would wish to economise on its use. Even where credit is free but collateralised there may still be an opportunity cost associated with the provision of collateral. In some countries, currently including the United Kingdom, central banks accept as collateral securities that banks (or at least most banks) need to hold for prudential requirements. In such circumstances, and if there is no remunerative way of using the securities intraday, for example because there is no intraday securities market, there would appear to be no opportunity cost. Banks, however, contest this. This may be because operational practicalities in using collateral intraday for payment purposes limit a bank's flexibility to trade the securities concerned (eg if substitution of securities is difficult or if the timing of the settlement of intraday repos prevents securities being used for an overnight transaction). Some banks' internal charging arrangements may also make collateral appear costly to payment managers, even where this does not reflect the true cost to the firm.

More significantly, there may be an opportunity cost for some banks in holding a portfolio that is different from the one that they would otherwise have chosen. This will certainly be the case in countries where there is no prudential requirement to hold the type and amount of securities used in payment systems. At the same time, central banks have an interest in maintaining the quality of the securities accepted as collateral to contain the risks they take. For all of these reasons, intraday credit may be costly for some banks. In any event, many banks seem to regard provision of collateral as being costly and want to minimise its use.

The desire to achieve the credit risk reducing properties of RTGS systems as well as the liquidity efficiency of DNS systems has encouraged the evolution of hybrid payment systems over the past couple of years.

Hybrid systems

To date two main types of hybrid system have emerged. The principal features of three such systems are summarised in Table 2 and further details are given in the Annex. One type is called continuous net settlement (CNS) and has evolved from DNS. Examples include the French and American systems, Paris Net Settlement (PNS) system and the New Clearing House Interbank Payment System (NewCHIPS), respectively. In such systems participants pay an amount of money into an account and all payment orders are entered at the beginning of the day into a queue. A computer algorithm searches for those that are 'largely' offsetting – ie they must not generate a net debit greater than the balance on the settlement account. If a group of payments is found that complies with this and certain other criteria, they are released and considered final. The algorithm runs frequently throughout the day and – in contrast with a DNS system – allows some, perhaps many, payments to be released and settled effectively in real-time.

Systems differ in how they settle any remaining payment orders. The PNS system, for example, requires participants to top up their account balances if the unsettled queue becomes too big. The intention is to

5: Canada's Large Value Transfer System uses a different approach to give participants an unconditional claim on the central bank. The Bank of Canada guarantees the settlement of LVTS, which in other respects is a protected DNS system.

6: The average per-minute overdraft is multiplied by the effective daily rate charged for overdrafts (currently 15 basis points) and then the value of the institution's deductible (derived from the level of overdraft that may be incurred without a fee) is subtracted to arrive at a daily charge. The daily charges are summed over two weeks and if they are less than $25, the fee is waived.

try to ensure that all payments are settled before the end of day. The NewCHIPS has an end of day process in which debtors should fund remaining net positions and it relies on the possibility of executing any payments that have not been settled in NewCHIPS through the parallel RTGS system (Fedwire).

Table 2: Features of three hybrid payment systems

Payemt system	NewCHIPS	PNS	RTGSplus[(a)]
Zero minimum balance	✓	✓	✓
Bilateral limits		✓	✓
Multilateral limits	✓[(b)]		✓[(c)]
Prefunding requirement	✓	✓	
FIFO settlement		✓[(d)]	
Bypass FIFO	✓		
Gross Settlement	✓	✓	✓
Bilateral offsetting			✓
Multilateral offsetting			✓
Bilateral netting	✓	✓	
Multilateral netting	✓	✓[(d)]	

(a) Reflects treatment of limit payments.
(b) Set at two times prefunding requirement.
(c) Set by payer to limit flow of liquidity to a group of payees.
(d) FIFO normally except in multiateral netting process.

The second type of hybrid system, illustrated by the RTGSplus system in Germany, seeks to incorporate a queuing facility into an RTGS system (a 'queue-augmented RTGS' system). Although some standard RTGS systems have queues to control the release of payments, the distinct feature of a queue-augmented system is that it looks to settle simultaneously payments that have been posted in the queue and that are broadly offsetting. The queue can work simply on a first-in-first-out (FIFO) basis where the priority given to payments depends only on when they arrive. A basic bilateral offsetting process looks only at mutual payments between a pair of participants. Alternatively, more sophisticated schemes can include an algorithm that searches the queue multilaterally for wider sets of payments, which are largely offsetting. Such an algorithm could also be written to override the time order of the FIFO system or to do so within categories of payments that have been accorded the same priority. This offsetting process involves simultaneous gross settlement, that is the individual payments are executed simultaneously in frequent batches. In legal terms settlement is gross (ie the individual obligations are not replaced by a net obligation) but it has the economic effect of netting payments because the offsetting gross payments are self-collateralising.

The two approaches would come to resemble each other the greater the proportion of transactions that are settled in the optimisation process – whether it involves netting or offsetting. The practical differences would be in the detailed design of the algorithm – how it selects which payments to settle. The choice of design will depend on the likely pattern of payments in the system and the needs of participants, for example, in terms of promptness of settlement. Both approaches could require similarly low levels of funds to achieve final settlement of all payments.

Both types of hybrid system need an associated RTGS system. In the case of a queue-augmented system, the system itself offers standard RTGS payment services, which may, for example, be used for high priority payments. In the case of the CNS system, an RTGS system will be necessary to provide the initial 'pump priming' provision of funds and to enable further payments to be made during the day or to execute uncompleted payments at the end of the day.

The choice between these alternative designs will depend on a number of factors, including the cost of building and operating them. For central banks, key factors will be how well the design contains settlement risk and whether the incentives the design creates for participants encourage a smooth flow of payments.

Settlement risk in payment systems

Settlement risk – the credit risk arising from the fact that payments made in the system may not be unconditional or irrevocable – is inherent in DNS systems. Only when settlement is complete can all payments become final (ie unconditional and irrevocable) between the direct participants in the system and the settlement risk ceases. A key aspect of DNS systems therefore is that the completion of an individual payment is dependent on the completion of all others during the same settlement period. That interdependence is at the root of the settlement risk in a DNS system. Only if all participants subsequently pay their net obligation to the system can settlement be assured. A protected DNS system seeks to mitigate these risks by providing procedures and resources to manage a settlement failure, but it will not guarantee settlement in all circumstances and the interdependence of payments remains.

CNS represents a significant advance on DNS in breaking the connection between the settlement of many payments. Payments that are completed in batches intraday can be settled with finality, and are not dependent on the subsequent settlement of other payments in the system. The link between the performance of settlement by participants at the end of the day and the settlement of a large group of payments – in practice, in the examples to date, a large majority of the value of a day's payments is settled before the end of the day – is broken. The less the amount left to be settled at the end of the day the less risky is the system, because less value remains at risk of an unwind. Some designs of CNS systems seek to ensure that all payments are settled with finality intraday and none are left to an end-of-day netting process.

RTGS systems, including those with queuing mechanisms, do not give rise to settlement risk. All payments are final as they are made, including those that are made simultaneously in batches in queue-augmented systems.

From a settlement risk perspective alone, there is a clear hierarchy of payment system design: unprotected DNS systems pose the greatest risk as a settlement failure could affect all payments, followed by protected DNS systems, CNS systems and RTGS systems that eliminate the possibility of a settlement failure. The differences between the systems, however, will be an empirical matter depending on the degree to which finality is achieved intraday. A CNS system, for example, may be capable of settling all payments without settlement risk if it completes all its payments in the matching mode and provided, importantly that the basis for the netting or similar process is legally robust.

The ability to settle payments safely – with little or no settlement risk – is a necessary but not a sufficient consideration in deciding whether a particular payment system design is best for a particular context. The way a payment system is used, the type of payments to be settled (in terms of the counterparties, volumes and values) and the incentives created by its design are also relevant considerations[7]. One illustration of the importance of the use of a system is the fact that a participant can face credit risk in a payment system even if the system itself (such as an RTGS system) is free of settlement risk. For example Bank A may have to make £100 millions of payments today as principal to Bank B, and Bank B may have an equivalent value to pay to Bank A. If Bank A pays first and Bank B fails before making its payment, Bank A bears the credit risk from the imbalance in the execution of payments. In principle an offsetting mechanism could help those banks manage their credit risk by making those payments simultaneously. In practice, many payments are made as agent rather than as principal and banks would not be able to offset the credit risk arising from the third party obligations. Nevertheless, banks still look at payment flows as an indicator of the financial and operational health of their counterparties, and the behaviour of banks in the payment system will depend on the incentives created by the design and operating rules.

The management of payment flows: incentives and behaviour

Banks typically want to manage their intraday liquidity carefully and thus want to see some balance in payment flows to ensure that the costs of providing liquidity to the system are fairly shared amongst the system members and to ensure that their own ability to make urgent payments is not inhibited. Appropriate design and operation of payment systems can help participants manage payment flows better and can provide incentives for collectively beneficial behaviour.

In a one-off 'game' in an RTGS model, for example, and assuming liquidity has a positive cost (whether a direct cost or an opportunity cost), the incentive on each participant would be to wait to be paid by others before making its own payments[8]. Collectively that would be sub-optimal because payments would tend to be made late in the day and there would be a risk that not all payments would be completed before the close of the system. Moreover, although participants would know what payments they had to make themselves, they would be unlikely to know the values of all payments due to them, particularly for their customers. Late payments and uncertainty about the total payment flows would complicate the management of their positions and of their treasury operations. Central banks also prefer to see payments completed reasonably early in the day because it reduces the risks of financial disturbances from

7: There are of course other relevant characteristics of good payment system design, such as low operational risk. For a detailed discussion of necessary characteristics see *Core Principles for Systemically Important Payment System*, BIS (2001).

8: This incentive is explored in Angelini (1998).

uncompleted payments and the risk that operational problems late in the day would disrupt the completion of payment business.

In practice, this 'game' is repeated daily and participants have an incentive to co-operate. One way of doing so is to develop rules of behaviour (whether explicit or implicit). In the United Kingdom, the banks have developed 'throughput' rules, which are agreed practices for paying each other promptly to avoid surprises and to encourage the smooth operation of payment systems. These are operated on an *ex post* basis where participants look to see whether their counterparties did make payments to the system as a whole sufficiently promptly (ie the rules operates on a multilateral basis)[9]. Such throughput rules assure participants that the initial funds held by participants in their accounts will be recirculated quickly intraday. There may also be conventions – implicit rules – where certain types of payments are carried out at particular times. In the US Fedwire system, for example, there are several short bursts of payment activity at predictable times where participants complete substantial volumes of payments[10]. This practice reduces the need for participants to draw on intraday credit from the Fed other than for very short periods, thus reducing their costs, and helps banks estimate their likely payment flows for the day.

These methods, however, are relatively imprecise and in practice many banks also operate *ex ante* controls on payment flows to individual counterparties. They set limits on the net payments they will make to individual counterparties and operate internal queues, monitoring flows in real-time and releasing payments steadily in response to incoming payments. As noted above, this is at least as much to conserve liquidity as to contain credit exposures.

The effectiveness of bilateral limits depends on the degree of concentration in the system and the extent to which payment flows are evenly distributed across participants. They can impose constraints by preventing the execution of 'circular' payment flows, for example when Bank A needs to pay Bank B, Bank B needs to pay Bank C, and Bank C needs to pay Bank A. In addition, depending on how they are set, there is the potential for asymmetric incentives. For example, a large bank may depend heavily on payments from another large bank and so have an incentive to establish prompt payment conventions with it, but the incentives to pay a small bank promptly may be less strong. To avoid this, banks need either a centralised monitoring function that can implement a more sophisticated multilateral throughput rule, ie one that ensures that no bank, big or small, 'free-rides' on the liquidity of others, or an automated approach such as a centralised queue.

The centralised queuing systems of CNS and augmented RTGS systems can help resolve these issues. A queuing system provides good incentives for participants to submit payments early because the greater the number and value of payments in the queue, the greater the probability of settling payments from other participants. Putting payments in the queue can release payments due to that participant from others. If a queue is sufficiently transparent so that a recipient can see intended payments, it will have a stronger incentive to enter payments of its own, and target more precisely the release of payments from other participants' queues. The linked payments do not, however, need to be perfectly offsetting as funds on the settlement account can be used as a buffer.

A centralised queuing procedure therefore has a number of potential advantages. First, the queues can be used to signal the willingness of a participant to pay its counterparties. In a bilateral system, a participant knows that if it is due to receive a payment in a second participant's queue, the payment will be released if it enters a similar payment order to the second participant. In a multilateral process more complicated payment chains may also be released even if there is no direct match of payments between two banks. Second, the queuing facility reduces the amount of liquidity needed to effect the payment. Neither party needs to find separate liquidity to support the payments. Simulations have shown that the effect can be substantial[11]. Third, it gives smaller banks access to this liquidity management tool if they do not have a sufficiently sophisticated in-house

9: The current UK rules are that 50 per cent of the value of payments must be made by 12.00 and 75 per cent by 14.30 on average over each month.

10: See McAndrews and Rajan (2000).

11: See Koponen and Soramaki (1998).

facility. This enables them to re-use liquidity quickly, sending it back into the system and speeding up the flow of payments for all.

Such approaches have a number of disadvantages. First, the cost of building and operating a more complex system may be high. Second, it may not yield substantial benefits if payment patterns are stable and payment conventions work or if there is a large number of participants and if payment volumes are widely spread. The chance of bilaterally-offsetting payments is much less in a large group, and multilateral algorithms are difficult to implement for systems with a large number of participants. Third, if there is a small number of large participants they may receive most of the benefit and smaller participants may need other ways (such as rules) of ensuring they share the benefits. Fourth, the complexity or uncertainty over the order and timing of settlement may concern some potential participants.

The signalling benefit of approaches using centralised queues is therefore likely to be most useful in a system in which there is significant day-to-day variation in individual banks' receipts. If daily flows are relatively stable in size, timing and counterparty there is less need for a bank to be able to signal its willingness to make payments at a particular time and standardised rules about payment timing may well be sufficient.

An issue for all types of hybrid systems is the status of messages in the queue. The queue indicates to the potential recipient the sender bank's provisional intention of remitting funds to the receiver. But if the receiver treats such information as equivalent to good funds, it can be taking the same risk as it would when crediting customers on the receipt of payment messages in a DNS system. For this reason, most queuing systems typically disclose only the amount of the payment and the beneficiary bank and release the customer details only when the payment is settled. There may, however, be a trade-off in the degree of transparency in the queues between giving sufficient detail to enable participants better to match and offset payments from counterparties, and avoiding the risk that participants treat queued messages as good funds, creating credit risk.

The UK system

In the UK, the domestic high-value payment system, CHAPS Sterling, is an RTGS system. The Bank of England provides intraday credit without charge against collateral, which is held for prudential purposes by the vast majority of participants. There is therefore seemingly little or no opportunity cost for most participants in obtaining credit. Participants nevertheless remain keen to manage daily liquidity carefully. To date the agreed multilateral throughput rules seem to have been broadly adequate in that payments flow promptly and smoothly. The increasing use of collateral intraday, for example, in supporting delivery versus payment in central bank money in CREST (the UK's securities settlement system), is increasing the attention given to liquidity management. The UK system has always had a gridlock-resolving mechanism called 'circles', which is similar to a multilateral offset algorithm, but it is rarely used, largely because CHAPS rules currently require banks to submit payments only if they have sufficient funds to settle them. The upgraded IT platform for the CHAPS system, called NewCHAPS, which went live in August 2001, has optional centralised queuing facilities available to banks that wish to avoid the costs of developing or upgrading their own in-house facilities. This system would be capable of incorporating offsetting payment techniques if the Bank and participants wished to introduce them.

Conclusions

There are many common elements driving the evolution of the design and operation of wholesale payment systems. Participants and central banks share the desire to reduce or eliminate settlement risk from payment systems by achieving prompt intraday finality, and to contain the amount of liquidity needed to support a payment system by improving the co-ordination of payment flows. The relative weights they give to these objectives may, however, differ. The range of possible solutions has been broadened by changes in technology that permit more sophisticated variants in payment system design to meet specific needs.

New designs involving centralised queues with efficient search algorithms can help achieve the safety and efficiency objectives when liquidity is costly. Both CNS and queue-augmented RTGS systems can enable participants to signal their willingness to pay counterparties if the counterparties are also prepared to pay them, and enable those linked payments to be executed simultaneously without the need for other, more costly, sources of liquidity.

The details of such systems, however, are important; for example the degree of transparency of the queues may affect credit risk. Such systems are most useful if there is substantial variability in the amount and timing of daily payment flows to different counterparties; existing techniques can reduce liquidity needs when payment flows are predictable. Multilateral approaches will maximise the liquidity efficiency for all, although they may be expensive or technologically difficult. Bilateral offsetting systems will work better where there is a small number of parties making most of the payments, and therefore a greater probability of matching payments, than where there are large numbers of participants and less concentrated payment flows. Such approaches, however, are likely to be needed only where there is a significant cost in the provision of intraday credit and where the values involved are substantial. In countries where one or both of these considerations do not apply, the operational and conceptual simplicity of a standard RTGS system might well mean it remains the preferred choice.

Bibliography

1: Angelini, P (1998) 'An Analysis of Competitive Externalities in Gross Settlement Systems', *Journal of Banking and Finance,* 22, pp. 1–18.

2: Bank for International Settlements (2001) 'Core Principles for Systemically Important Payment Systems', January (www.bis.org/publ/cpss43.htm).

3: Bank for International Settlements (1997) 'Real-time Gross Settlement Systems', March (www.bis.org/publ/cpss22.htm).

4: Bank for International Settlements (1990) 'The Report of the Committee on Interbank Netting Schemes of the Central Banks of the Group of Ten Countries' (Lamfalussy Report), November (www.bis.org/publ/cpss04.htm).

5: Bank of England (1989) 'Challenges Facing the Sterling Wholesale Payment Systems', *Bank of England Quarterly Bulletin*, August, pp.401–406.

6: Beau, D, Ducher, J and Ryckebusch, G (2000) 'Operating Characteristics of France's TBF and PNS High-Value Payment Systems: Lessons from Modeling', Payment Systems Worldwide, No 1 Vol 11, pp.20–29, 48.

7: Board of Governors of the Federal Reserve System (1984) 'Measures to Reduce Risk in Large Electronic Fund Transfers', *Federal Reserve Bulletin*, April, pp.329–330.

8: Clearing House Interbank Payment System Company (2001) 'CHIPS Rules and Operating Procedures', January.

9: Deutsche Bundesbank (2000) 'RTGSplus – the Bundesbank's New System for Individual Payments', *Deutsche Bundesbank Monthly Report*, June, pp.59–72.

10: Deutsche Bundesbank (2000) 'RTGSplus: The Modern Liquidity Saving Real-Time Gross Settlement System for Euro', July.

11: Deutsche Bundesbank (1999) 'Convergence of Gross and Net Systems – the EAF/ELS Example', September, SIBOS, Munich.

12: Fry, J, Kilato, I, Roger, S, Senderowicz, K, Sheppard, D, Solis, F and Trundle, J (1999) 'Payment Systems in Global Perspective', Routledge International Studies in Money and Banking.

13: Koponen, R and Soramaki, K (1998) 'Intraday Liquidity Needs in a Modern Interbank Payment System: a Simulation Approach', Bank of Finland Studies, E:14, Helsinki.

14: McAndrews, J and Rajan, S (2000) 'The Timing and Funding of Fedwire Funds Transfers', *FRBNY Economic Policy Review*, July.

15: The Payments Risk Committee (2000) 'Intraday Liquidity Management in the Evolving Payment System: a Study of the Impact of the Euro, CLS Bank, and CHIPS Finality', April, Intraday Liquidity Management Task Force, New York.

Annex: Three examples of hybrid systems

France – PNS

PNS is a CNS system that provides real-time irrevocable settlement. It has been operational since 19 April 1999. It evolved out of Systeme Net Protégé (SNP), a protected DNS system secured by collateral, which operated within bilateral and multilateral limits.

PNS operates by settling positions in central bank money throughout the day using netting or optimisation procedures and it is intended for high-value payments less urgent than those handled in Transferts Banque de France (TBF, the French RTGS system). Liquidity may be moved between TBF and PNS at any time but each participant must transfer a minimum of €15 millions to PNS each morning. Banks whose overall net position for the day is negative are required to bring in the necessary liquidity progressively during the day to settle all queued payments, in order that no payments are rejected at the end of the day.

Participants are able to control flows to counterparties by setting bilateral limits. In addition, each participant's settled position vis-à-vis the system as a whole must remain positive. As no bank can be a net debtor to the system as a whole, no default fund is necessary. PNS payments will be settled if they are within bilateral limits and sufficient liquidity is available. Payments that would breech a bilateral limit or which would cause a participant's account balance to become negative are queued. If there is an existing queue, lower value payments (below €1 million) will be given priority to avoid overloading the deferred settlement process. Queued payments may be settled by one of three processes: bilateral optimisation, queue scanning or multilateral partial optimisation. The first two are continuous and follow the FIFO principle whereas the third departs from FIFO to settle as many payments as possible. The bilateral limits apply whichever process is used.

Bilateral optimisation is triggered whenever a payment is added to the queue. The system examines all the queued transactions between the paying participant and each other participant, and attempts to process the largest number of transactions simultaneously, consistent with FIFO and the bilateral limits. Queue scanning occurs after each settled payment and each successful bilateral optimisation. Once an account has been credited, the system checks whether any queued transactions on that account can then be settled.

Multilateral partial optimisation is carried out whenever the system operator notices gridlock. This mechanism involves a number of passes over the data, each pass resulting in some payments being dropped. First, individual payments, which would cause bilateral positions to exceed bilateral limits, are dropped. Second, the queue of the participant with the largest virtual debit balance is examined and transactions are removed until the participant's position becomes positive. If there are several transactions that could be removed to achieve this, then the most recent are removed. This second operation is repeated until no debit balances remain and the remaining transfers are settled simultaneously.

Germany – RTGSplus

RTGSplus was introduced on 5 November 2001. It replaced Germany's two large value payment systems, Euro Link System (ELS) and Euro Access Frankfurt (EAF), incorporating features of both. ELS was an RTGS system and the German component of TARGET (Trans-european Automated Real-time Gross settlement Express Transfer system), while EAF was a CNS system in which payments were settled through a matching mechanism supported by some pre-funded balances. RTGSplus processes the payments of both systems and offers participants a number of controls to manage their payments.

RTGSplus payments are made against balances held on dedicated settlement accounts. Liquidity may be transferred through TARGET to and from an RTGSplus account throughout the day. Participants without access to TARGET can use correspondent arrangements. Surplus funds are returned to a designated account at the end of the day.

RTGSplus payments can be either express or limit payments. Express payments are individual RTGS payments, settled on a FIFO basis, using any of the liquidity available to the paying bank. Limit payments are settled on a bypass-FIFO basis and may access only a part of the liquidity on the settlement account.

Banks can manage their payments by choosing which payments to put through the express route and reserving liquidity for those payments, by setting bilateral and multilateral limits for limit payments, and by reordering or revoking payments. Multilateral limits apply to counterparties for which there is no bilateral limit in place. Participants can monitor both the balance in their settlement account and detailed information on incoming and outgoing payment queues.

Incoming payments will be settled in real time if there are sufficient funds and any limits in place are met. If a payment fails to settle, it will be queued. The two types of payments are queued separately and the express payments queue takes priority. Queued express payments are checked to establish whether they can be settled whenever the settlement account is credited, when the order of payments is changed, on revoking an express payment or on changing the payment type. If a participant's express payments are queued, its limit payments cannot settle. An algorithm runs continuously to find queued limit payments that can settle. It searches for bilateral or multilateral offsets first using FIFO but then bypassing FIFO to allow subsequent payments to settle ahead of queued payments to maximise settlement subject to any limits in place. Any payments still in the queue at the end of the day are rejected.

US – NewCHIPS

CHIPS was converted to a CNS system on 22 January 2001. It continuously matches, nets and settles queued payments. Previously CHIPS was a protected DNS system operating with net sender debit caps and net bilateral credit limits. It also had a collateral pool, which could accommodate the end of day failure of the two largest net debtors.

Under the new system, at the start of the day, participants must pre-fund their account with an amount called the 'initial balance requirement'. This requirement is set using a formula based on recent transaction history. It is recalculated weekly. Participants prefund their CHIPS settlement account at the Federal Reserve Bank of New York through Fedwire. If the participant is not a member of Fedwire it must have arrangements with a member to provide the initial balance requirement. Funds cannot be added or withdrawn during the day and surplus funds are returned through Fedwire at the end of the CHIPS business day.

Priority payments may be flagged to be considered first for release. It is also possible to delete unreleased payments from the queue until the system closes for the receipt of new payments (at 17.00 Eastern Time). Constraints are placed on participants' intraday positions such that the minimum position during the processing day cannot be less than zero and the maximum position cannot be more than twice the initial balance requirement.

Payments may be settled individually, netted bilaterally or netted multilaterally. The algorithm searches the queues for individual payments, or batches, whose settlement can be accommodated within the constraints, including that of the available balances. Although consideration is given to the order in which the payments arrived, the algorithm selects payments to be released on a best-fit basis. It will bypass payments in the queue that are not likely to be part of a batch that can be released.

Once the system is closed for new payments, the maximum position control is removed and the algorithm is run to match, net and set off as many of the remaining unsettled payments as possible. A closing position is then calculated for each participant comprising the multilateral net position for the remaining payments plus its current balance. Banks with negative closing positions are required to fund their 'final prefunding requirement' via Fedwire within 30 minutes of the notification of this position. Once all the transfers are made, all remaining payments are released and settled and CHIPS sends credits via Fedwire to banks with a positive closing position. If one or more of the debtor banks do not provide their final pre-funding requirement in full, the algorithm runs to release and settle as many of the remaining payments as possible. Any payments still remaining in the queue are rejected by the system, the affected sending banks are notified, and CHIPS credits banks with a positive closing position via Fedwire. The unsettled payments would normally be redirected by the banks via correspondents or through Fedwire that business day.