BANK OF ENGLAND

Working Paper No. 355
The network topology of CHAPS Sterling

Christopher Becher, Stephen Millard and

Kimmo Soramäki

November 2008

BANK OF ENGLAND

Working Paper No. 355

The network topology of CHAPS Sterling

Christopher Becher,(1) Stephen Millard(2) and Kimmo Soramäki(3)

Abstract

In this paper, we seek to understand the network topology of large-value interbank payment flows in the United Kingdom so as to understand better the risks associated with the system. We first examined the broad network topology of interbank payments in the United Kingdom. We found that, despite the fact that there are far fewer banks in the United Kingdom than in the United States, the structure of UK interbank payments is similar in certain respects to that of the United States, but that the tiered structure of the UK system implies rather different risk characteristics. We then looked at CHAPS and found that payment flows in CHAPS form a well-connected network whose properties change little day to day. This means that liquidity is able to flow efficiently around the network and that the network is quite resilient to shocks. This finding was backed up by examining the effects of a particular incident on the properties of the CHAPS network. In that particular instance, the effective removal of one bank for much of the day had little impact on the ability of other banks to make payments between one another.

Key words: Network topology, CHAPS Sterling, interbank payment flows.

JEL classification: D85, G21.

---

(1) European Commission. Email: chrisbecher@yahoo.com
(2) Bank of England. Email: stephen.millard@bankofengland.co.uk
(3) Email: kimmo@soramaki.net

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or of the European Commission. The paper has benefited from comments from an anonymous referee, Stephen Denby, Marco Galbiati, David McFarlane, Mark Manning, Andrew Mason, Ben Norman and Gabriel Sterne as well as seminar participants at the Bank of England. This paper was completed while the first author was at the Bank of England. This paper was finalised on 7 August 2008.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

One of the core purposes of modern central banks is to contribute to financial stability. This entails assessing risks across the financial system as a whole – systemic risk – that would otherwise undermine the system in general, and seeking to make the system stronger by reducing such risks. Payment systems, which facilitate transactions between individuals, businesses and financial institutions, form a crucial part of the financial system and play a vital role in ensuring the smooth implementation of monetary policy. So, it is important for central banks to understand how shocks in one institution can be propagated across payment systems if they are to seek to reduce the systemic risk in such systems.

One way of trying to do this is to characterise the structure of a payment system – its 'network topology' – using tools recently developed by physicists. Once we understand the structure of the network of banks and the payments they make to/receive from each other, we can assess the stability of this network to particular shocks. In this paper, we seek to understand the network topology of large-value interbank payment flows in the United Kingdom so as to understand better the risks associated with the system.

The UK large-value payment system – the Clearing House Automated Payment System (CHAPS) – consists of only 15 banks. Banks that are not direct members of the system (so-called 'second-tier' banks) have to make their payments via a correspondent bank that is a member of the system. We first examine the ramifications of this tiered structure, and illustrate the broad network topology of interbank payments in the United Kingdom, using data from the 2003 CHAPS Traffic Survey. We find that, despite the fact that there are far fewer banks in the United Kingdom than in the United States, the structure of UK interbank payments is similar in certain respects to that of Fedwire (the US large-value payment system). But while the two networks are in some respects similar, the tiered structure of the UK system implies rather different risk characteristics.

We then look at the CHAPS system as a network containing only the settlement banks. We find that payment flows in CHAPS form a well-connected network – every bank is connected to every other bank by some set of payment flows – and that its properties change little day to day. A consequence of this network structure is that liquidity is able to flow efficiently around the network. We also find that the network develops only gradually in the early hours of opening. The explanation for this pattern lies in the purposes of payments being made at this time, and in particular the tendency to withhold payments until time-critical payments have been settled. We also saw slight peaks in the number of pairs of banks involved in payments before noon and late afternoon, indicating that the network is particularly busy at these times. This variation indicates that the impact of an operational disruption may vary according to the time of day at which it strikes.

Finally, we examine the effects of a particular incident – where one of the banks was unable to make payments for a large proportion of the day – on the properties of the CHAPS network. The network appears to be highly resilient to this type of shock. In the particular instance of an operational outage examined here, the effective removal of a node for much of the day had little

impact on the ability of other banks to make payments between one another. The fact that the network is ‘well-connected’ will have contributed to this resilience. However, we cannot discount the possibility that operational disruptions at one or more large banks, especially if they are net suppliers of liquidity to the system, would have a more severe impact on the payment network than was evident from this case study.

1 Introduction

One of the core purposes of modern central banks is to contribute to financial stability. This entails assessing risks across the financial system as a whole – systemic risk – that would otherwise undermine the system in general, and seeking to make the system stronger by reducing such risks. Payment systems, which facilitate transactions between individuals, businesses and financial institutions, form a crucial part of the financial system and play a vital role in ensuring the smooth implementation of monetary policy. So, it is important for central banks to understand how shocks in one institution can be propagated across payment systems if they are to seek to reduce the systemic risk in such systems.

Recently, physicists have made progress in understanding the structure and functioning of complex networks, of which payment systems can be thought to be examples. This literature has sought to characterise the structure of complex networks and assess the stability of such networks to particular shocks.(1) Boss *et al* (2003) have brought these techniques into the field of financial stability by examining the network topology of interbank exposures in Austria with a view to assessing the amount of systemic risk in the Austrian banking system. Similarly, Soramäki *et al* (2007) have examined the network topology of interbank payment flows through Fedwire in the United States and analysed the effects of the 11 September 2001 terrorist attacks within this framework. Inaoka *et al* (2004) look at networks formed from payments in the Japanese BoJ-Net system and Lublóy (2006) investigates the topology of the Hungarian interbank payment system.

In this paper, we seek to understand the network topology of large-value interbank payment flows in the United Kingdom. In doing this, we seek to use this knowledge to understand better the risks associated with the system as it is now; we do not, in this paper, say anything about the form a welfare-maximising network should take. One aspect of the topology of UK large-value interbank payment flows is immediately apparent. The UK large-value payment system – the Clearing House Automated Payment System (CHAPS) – consists of only 15 banks.(2) Banks that are not direct members of the system (so-called ‘second-tier’ banks) have to make their payments via a correspondent bank that is a member of the system. In the paper, we examine the ramifications of this tiered structure and illustrate the broad network topology of interbank payments in the United Kingdom. These results are presented in Section 3.

But, in understanding the impact of an operational failure affecting one of the settlement banks on CHAPS payments, it is probably most instructive to look at the CHAPS system as a network containing only the settlement banks. We do this in Section 4. We find that payment flows in CHAPS form a well-connected network and that its properties change little day to day. What is

---

(1) See, for example, Albert *et al* (1999) and Crucitti *et al* (2004).

(2) The Royal Bank of Scotland (RBS) and National Westminster (NatWest) Bank currently retain and use separate settlement accounts in CHAPS; however, for the purposes of this study, we combine the two banks into a single node. The other 14 banks are Abbey, ABN Amro Bank, the Bank of England, the Bank of Scotland (HBOS), Barclays Bank, Citibank, Clydesdale Bank, CLS Bank, the Co-operative Bank, Deutsche Bank, HSBC Bank, Lloyds TSB Bank, Standard Chartered Bank and UBS. At the t.ime this research was carried out, UBS were not a direct member of CHAPS Sterling and so, when we consider the network of the settlement members, we only have a network consisting of 14 banks. All references to ‘CHAPS’ in this paper refer to CHAPS Sterling only.

more interesting is how the properties of the network change over the course of a day, since knowing this will enable us to assess the stability of the system to operational failures of different members at particular times of the day. So, in Section 5, we examine how the network properties change over the course of the day. In Section 6, we examine the effects of a particular incident – where one of the banks was unable to make payments for a large proportion of the day – on the properties of the CHAPS network. We use this example as a way of assessing the robustness of the network to this type of shock. Finally, Section 7 concludes.

2 Definitions

Drawing on the network topology literature, we make use of a number of key terms throughout this paper. A *network* consists of a set of *nodes* connected by *links*, which can be directed – ie, crossed in one direction only – or undirected. The *weight* of a link indicates the importance of that link and the *degree* of a node refers to the number of links that originate (*out degree*) or terminate (*in degree*) at that node. A sequence of nodes in which each node is linked to the next is termed a *walk*; if all nodes in the walk are distinct, the walk is a *path*. The *distance* between two nodes is the length of the shortest path between the nodes.

The properties of networks can be compared using these basic features, for example by calculating the average path length between nodes, or the average degree of a node in the network. More complex measures can also be derived. The *connectivity* of a network, for example, is the unconditional probability that two nodes in a network share a link, or, in other words, the number of actual links as a proportion of the number of potential links. And the extent of *clustering* can be measured as the probability that two nodes which neighbour another node themselves share a link, measured as the number of directed links between the neighbours of a node over the number of potential links between them.

A network in which all nodes have a link to all other nodes is a *complete* network. A *component* of a network is a subset of nodes within which any two nodes can be connected by a path. If a network consists of a single component, it is a *connected* network; if more than one, it is a *disconnected* network (since not all nodes can be reached by every other node). The largest component in a network in which all nodes are connected to each other via undirected paths is termed the *Giant Weakly Connected Component* (GWCC). A *Giant Strongly Connected Component* (GSCC) comprises all nodes that are connected through a directed path.

In modelling CHAPS payment flows as a directed network, banks are represented as nodes in the network and payments between banks form the links between these nodes. We can think of these links as being ‘directed’: if *A* only makes (but does not receive) payments to *B* then there would be a directed link from *A* to *B* but not one from *B* to *A*; if both banks made payments to each other, we have two directed links, one in each direction. The weight attached to a link is proportional to the value or volume of payments passing through that link. The design of CHAPS is such that all members could technically make payments to all other members, hence CHAPS payments could in principle be modelled as a complete network. The empirical question is to what extent each of these links is used in practice, and the implications of this for the flow of liquidity around the network.

3 The topology of CHAPS payment flows: all banks

As mentioned in the introduction, CHAPS is a highly tiered system, in which a large number of indirect participants make payments using agency agreements with a small number of settlement banks. Underlying the relatively simple network of payments between the direct members, therefore, there is a much more extensive network of payment flows between the banks that originate payments and those that ultimately receive them. This situation contrasts with that of eg, Fedwire, in which many more financial institutions access the clearings directly, although smaller institutions do use correspondent banking relationships to effect payments. In light of this, it is more instructive to compare the characteristics of Fedwire with the network of CHAPS payments made by both direct and indirect members of the system, rather than with the settlement bank network.

To examine CHAPS payments made by both direct and indirect members, we use data from the 2003 CHAPS Traffic Survey. The survey consists of a sample of CHAPS payments for five days in February 2003. For each of the payments in the sample, the data contains the value and purpose of the payment, when (date and time) it was sent, and, importantly for this paper, codes identifying the bank sending the payment and the bank receiving the payment (whether or not they are actually direct members of CHAPS) together with codes identifying the direct members of the system between which the payment was actually executed.(3)

There are two approaches to visualising the network of large-value sterling payment flows. The first is to consider only the payment relationships between the ultimate payer and payee bank.(4) Alternatively, given that these payments are made via the payer and payee settlement banks, each payment could be viewed as consisting of up to three separate flows: from the payer to the payer's bank; from the payer's bank to the payee's bank; and from the payee's bank to the payee.(5) Results based on both of these approaches – which we term the 'relationship' and 'flow' approaches, respectively – are presented in Charts 1a and 1b.

The network of settlement banks is clearly visible at the core of the tiered network, even when viewed using the 'relationship' approach. (The characteristics of this network are examined explicitly in Section 4.) There is wide variation in node and link strength within the tiered network. Indeed, some of the 'second-tier' banks originate or receive payments to a value similar to those of some of the settlement banks, as shown by the larger white nodes in Chart 1a. This may suggest that the perceived cost of becoming a direct member of CHAPS is relatively high for these banks, although it may be the case that these banks had unusually high payment flows on the dates when the traffic survey was carried out and that in normal times their payment flows are not as large. But equally, the fact that large payments are made on behalf of second-tier banks creates large intraday exposures between these banks and their settlement

---

(3) We are extremely grateful to the CHAPS Clearing Company Ltd. for allowing us access to these data.

(4) The network of flows between banks is in fact underpinned by another network of flows between their customers, including households and corporates. We are not able to identify the ultimate originator and recipient of a payment in our data set.

(5) We should note that, if the ultimate payer and payee of a particular payment hold accounts at the same bank, this payment will be 'internalised' (ie, will not go through the CHAPS system) and so not be captured in our data.

banks.[6] The flow approach suggests that some indirect participants may route payments through more than one settlement bank (Chart 1b). This suggests that they try to diversify their exposure to credit and operational risk across settlement banks.

Chart 1: Value-weighted topology of 'tiered network', 2003

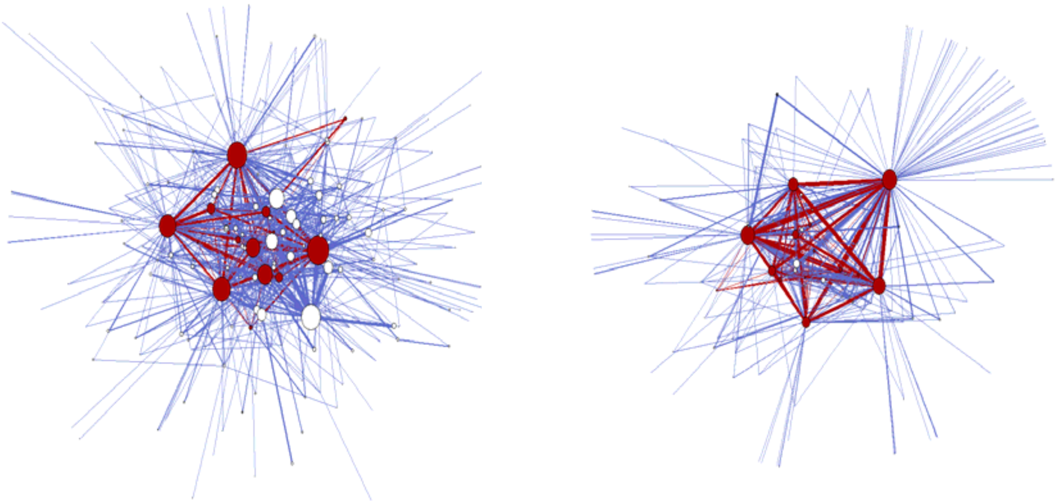

a) Relationships **b) Flows**

Notes: CHAPS settlement banks and the links between them are displayed in red; non-settlement banks are displayed in white and the links among them, and between them and the settlement banks, are shown in blue. The diameter of each node is proportional to the square root of the value of payments sent and received by the bank; and the weight of the link (thickness of the line) is proportional to the value of payments passing through that link.

Since all nodes in the tiered network are connected to each other, the entire network can be characterised as a GWCC. Within this, a GSCC can be identified, in which all nodes are connected via directed paths (ie, sets of payments flowing in one direction). To illustrate exactly what we mean by this in a payments context consider a network of only four banks: *A*, *B*, *C* and *D*. Suppose *A* and *B* make payments to/from each other. *A* has no dealings with banks *C or D*. *B* makes payments to *C* (as well as *A*) but not *D* and does not receive any from *C* or *D*. Finally, *D* only makes payments to *C* and does not receive any payments. In this case, the GWCC will consist of all four banks, since they are all connected to each other. The GSCC will consist only of banks *A*, *B* and *C* since payments flow from *A* to *B* to *C* in one direction but only from *D* to *C* and *B* to *A* in the other direction.

The properties of the components of the CHAPS tiered network are summarised in Table A, alongside results from a similar study of Fedwire.[7] Clearly, the Fedwire network is the larger by a considerable margin, with an average (GSCC) network size of 5,086 compared with around 120 direct or indirect participants in the CHAPS GSCC (and only 14 banks in the network of settlement banks). It is thus to be expected that the number of links and the average degree is much higher in Fedwire. Consistent with the more concentrated nature of the network, connectivity is higher in the CHAPS GSCC (5.1% compared to 0.3% in Fedwire). However, despite the smaller number of nodes, the average path length in the CHAPS GSCC in 2003 was

---

[6] Harrison *et al* (2005)
[7] Soramäki *et al* (2007)

2.6 and 2.4 in the GWCC, similar to the average length observed in Fedwire (2.6). The degree of clustering, measured by the clustering coefficient, is lower in CHAPS than in Fedwire (0.27 in the CHAPS GSCC, compared with 0.53 in Fedwire).

Table A: Properties of the CHAPS tiered network

	CHAPS GWCC	CHAPS GSCC	Fedwire GSCC
Number of nodes	337	117	5,086
Number of links	989	692	76,614
Connectivity (per cent)	0.9	5.1	0.3
Average degree	2.9	5.9	15.2
Maximum out degree	52	48	1,922
Average path length	2.4	2.6	2.6
Clustering coefficient	0.23	0.27	0.53

But while the two networks are in some respects similar, the tiered structure of the CHAPS system implies rather different risk characteristics. Correspondent banking relationships are used to make payments in both systems but are more prevalent in the case of CHAPS: the CHAPS Traffic Survey, 2003, suggests that around 30% of CHAPS payments originate from correspondent banks, which include many other domestic banks. Tiering potentially introduces risks to the system, in particular as a result of credit exposures between first and second-tier banks, and the concentration of all payment activity at a small number of settlement banks.(8) These risks would be mitigated were some of the large 'second-tier' banks illustrated in Chart 1 to pursue direct membership of the system. However, this would entail additional operational costs for these banks. There are also potential benefits from a highly tiered system. Liquidity pooling and the internalisation of payments may reduce liquidity demands. Moreover, the relatively small direct membership may serve to facilitate co-ordination between banks, which in turn may help to improve the efficiency of liquidity recycling in the system.

4 The topology of UK payment flows: settlement banks

We next consider the network created by the payment flows between the CHAPS settlement banks, looking first at the properties of the network on a single day. We then assess how stable is the network topology over time by examining the evolution of some of these measures over the period from July 2005 to June 2006. For this we use the Bank of England's own 'Payments Database'. This database has a record of every transaction that is carried out over the CHAPS Sterling system. For each payment the database records the value of the payment, when (date and time) it was sent, and the sending and receiving settlement banks: that is, the two direct members of the system involved in processing the payment but not the ultimate sender and receiver of the payment.(9)

---

(8) Harrison *et al* (2005) discusses these risks in more detail.

(9) As the operator and settlement agent of the CHAPS Sterling system, the Bank is involved in all CHAPS Sterling transactions and maintains a record of them for research purposes. James (2003) describes the data in more detail.

Chart 2 provides a visualisation of the CHAPS network on a sample day.(10) The thickness of the links is proportional to their weight, defined here as the value of payments passing through the link. It is clear that payments between settlement banks form a well-connected network, consisting of a single component.

The high level of connectivity is confirmed by the descriptive statistics presented in Table B. The network displays both very high connectivity (88%) and a short average path length (1.1).(11) Indeed, the network is almost complete: most members have directed links with most other settlement banks and – with the exception of links with CLS Bank – all links are bi-directional.(12) The average degree of a node is 11.4; that is, an average of just over eleven links originate from each node.

Chart 2: Value-weighted topology of settlement bank network (17 May 2006)

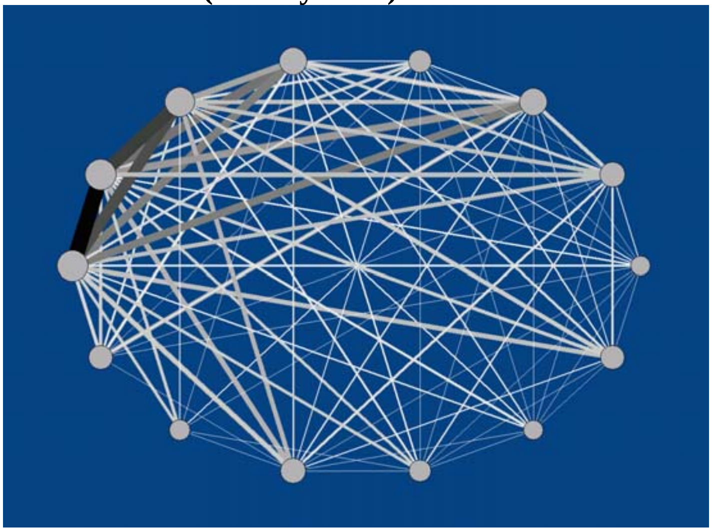

Table B: Properties of the CHAPS 'settlement bank network' (17 May 2006)

Nodes	14
Links	160
Total value out (£ billion)	188
Mean value out per node (£ billion)	13.5
Total volume out	105,938
Mean volume out per node	7,567
Connectivity (per cent)	88
Maximum / average / minimum out degree	13 / 11.4 / 6
Average path length	1.1
Average clustering coefficient	0.92

---

(10) Although we only consider one particular day, our results later in the section suggest that the network properties of CHAPS Sterling are quite stable day to day and so the results presented here are unlikely to be atypical.

(11) Note that connectivity refers to the proportion of links actually used on the sample day. All of the links are in principle capable of being activated on any given day.

(12) We discuss the particular characteristics of CLS that cause this below.

Not only are payments concentrated among a small number of settlement banks but there is also, as Chart 2 indicates, a high degree of concentration among a subset of these settlement banks (clustered on the left-hand side of the chart). Indeed, four banks account for around 80% of payments by both value and volume, as illustrated by Chart 3. This suggests that the impact of an operational disruption – or credit event – at certain critical nodes, is potentially high in CHAPS; we return to this issue in Section 6.

Even members that account for a comparatively low value of CHAPS payments are highly connected, as evidenced by the wide dispersion of the strength of high-degree nodes shown in Chart 4. Of the 12% of potential links unused on the sample day, over half are accounted for by the low degree of just two banks, who joined the system in late 2005. As one bank, Abbey, was still relatively new to the system when we carried out the analysis, we might expect its node degree to have increased since then. Indeed, the links that were not used on the sample day were used on other days. The other, CLS Bank, is a 'special-purpose' bank and, as such, will behave differently to the other CHAPS members.

Chart 3: Concentration of CHAPS payments

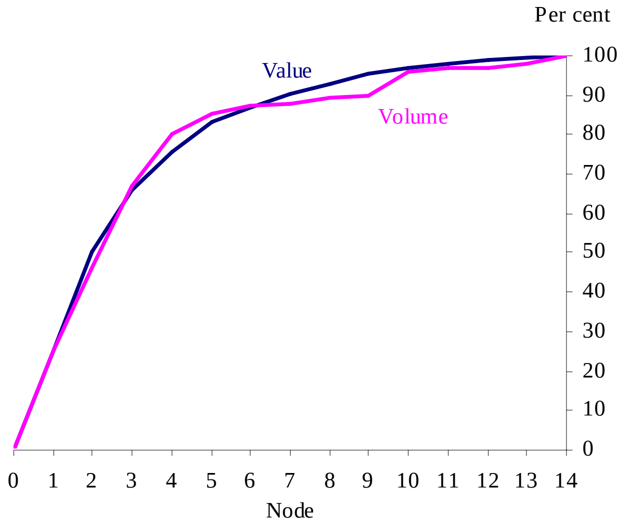

Chart 4: Node strength and value out

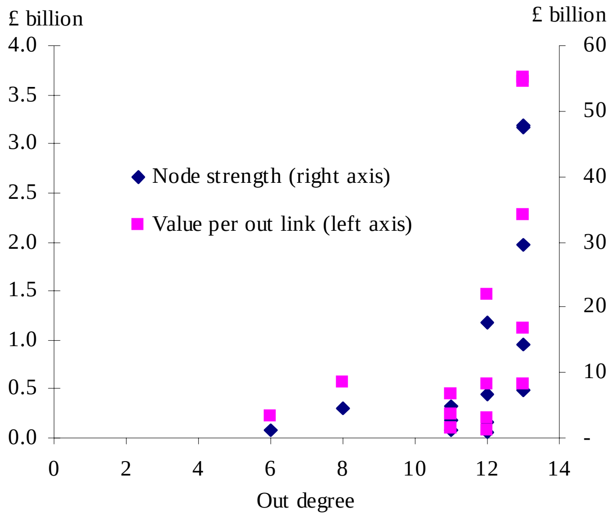

CLS Bank acts as a settlement bank for foreign exchange transactions, enabling such transactions to be carried out on a payment-versus-payment basis. For each member, CLS calculates the member's net position in each currency in which it trades. If a member has a credit position with CLS in a particular currency, then CLS will make payments to that member over the relevant large-value payment system in that currency; if a member has a debit position in a particular currency, then it will make payments to CLS over the large-value payment system in that currency.

The upshot of this for CLS payments in CHAPS is that CLS Bank will only make/receive payments to/from its settlement members; four CHAPS settlement banks are not currently members of CLS. A second implication is that those banks with net short positions in sterling on a given day will only pay *in* to CLS in the CHAPS system on that day; and only those with net long positions in sterling will receive a CHAPS payment from CLS. So, we might expect to find the links between CLS and the settlement banks to be one-way (if there are links at all).

However, this will not *always* be the case since some banks act as nostro agents for third parties; it is possible that they will need to make payments into CLS, say, while receiving payments from CLS on behalf of these third parties. We can also note that the directions and weights of the links between commercial banks and CLS Bank will, thus, vary from day to day.

Turning now to the evolution of the CHAPS network over time, Charts 5-8 illustrate that the characteristics of the CHAPS payments network remained reasonably stable after the introduction of the two new members in late 2005. Prior to this time, CHAPS payments formed a complete, fully connected network. Following the introduction of the new members, the average degree naturally increased, but also became more volatile (11.4 with a standard deviation of 0.16, compared with 11.0 with a standard deviation of 0.10 before), while the connectivity of the network fell from unity to an average of 0.88 (with a standard deviation of only 0.01) in two steps: the first on 14 November 2005 when Abbey joined and the second on 28 November 2005 when CLS joined. The increase in volatility relates in large part to volatility in the connectivity of the new members – particularly that of CLS for reasons explained earlier – rather than volatility in the characteristics of the existing network. There were no subsequent step changes in the values of these statistics.

Chart 5: Out degree

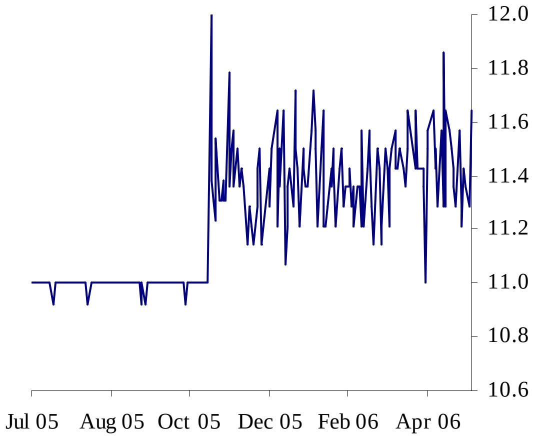

Chart 6: Average connectivity

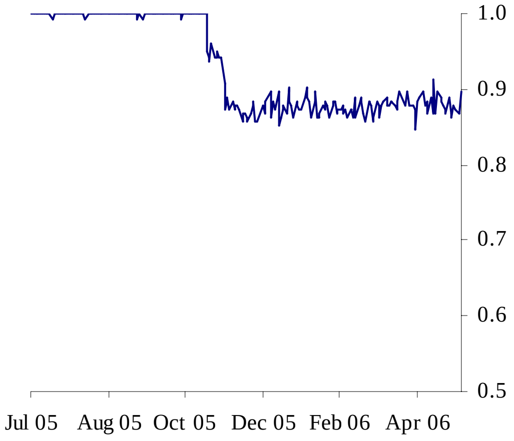

Chart 7: Average path length

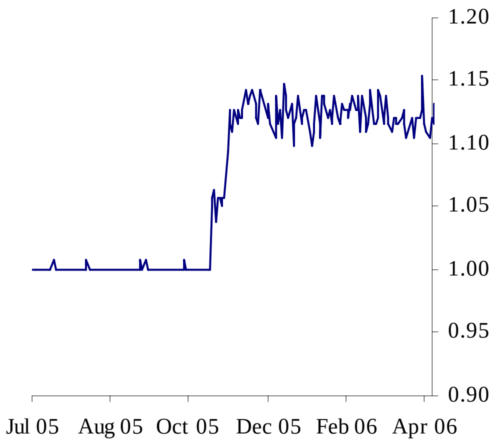

Chart 8: Average clustering co-efficient

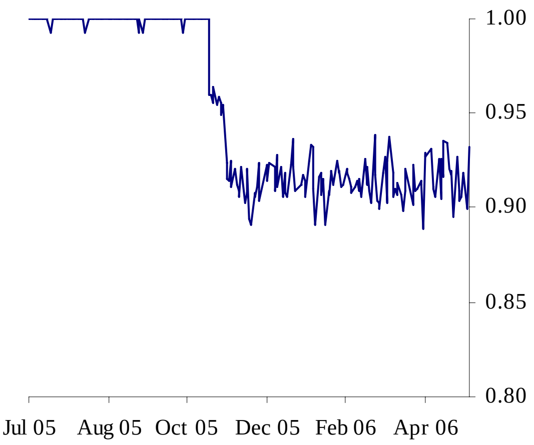

5 The intraday evolution of the CHAPS network

Hitherto, we have discussed the properties of the network of CHAPS payments resulting from the accumulation of all payment flows over the course of the day. However, the network will not exhibit these properties at all times; for example, we might expect to see payments within the network building up over several hours with the result, for example, that the level of connectivity is lower at the beginning of the day. These intraday characteristics have important implications for the impact of an operational disruption affecting a member at particular times of the day, and so the identification of these characteristics can contribute to the design of tools to mitigate these risks.

To consider these effects, we now analyse the intraday development of the network, using two methodological approaches. First, we allow links within the network to accumulate and then calculate the network properties every 20 minutes. Second, we treat each 20-minute interval as a separate network and measure the network characteristics within each period, ie, we take a snapshot of the network based on payments made between 5.40 am and 6.00 am; then of the network based on payments made between 6.00 am and 6.20 am; and so on. So, for example, the 'cumulative network' at 10.00 am is the network resulting from all payments made up to that point, whereas the 'snapshot network' is constructed solely from payments made between 9.40 am and 10.00 am. In each case, the data is drawn from four randomly selected days between July 2005 and June 2006.

The cumulative results illustrate how the network characteristics described in Section 4 evolve over the course of the day. As Chart 9 shows, peak connectivity (of around 88%) is not achieved until the early afternoon. At 7.00 am, after one hour of payment activity, connectivity is just over 50%, reaching 80% by 9.00 am. The 'snapshot' statistics reveal that connectivity during any given 20-minute interval is (unsurprisingly) lower than that of the cumulative network; on average, only around half of the available links are used during any 20-minute interval. The network is most active from late morning onwards, with slight peaks before noon and in the late afternoon.

Chart 9: Intraday networks, connectivity

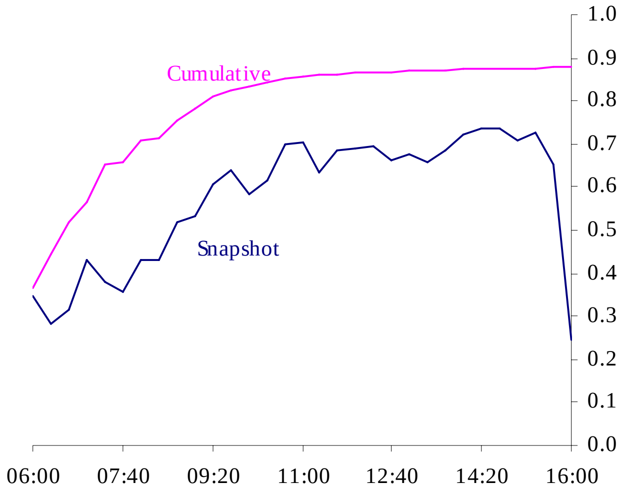

Chart 10: Intraday networks, out degree

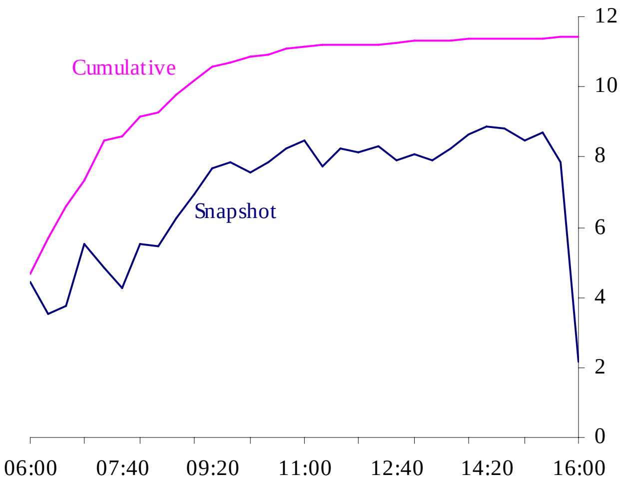

We observe the same pattern in the evolution of the average degree of the nodes in the system. Average node degree in the cumulative network reaches seven by 7.00 am (that is, each bank is connected to an average of seven other banks) and rises to ten shortly before 9.00 am. The maximum (11.4) is not achieved until mid-afternoon. The average degree within the 20-minute snapshot networks remains close to eight for much of the day, reaching a maximum of just under nine late in the day.(13)

These patterns in network activity correlate well with established features of the payment day. In particular, the tendency of the network to become busier in the late morning is consistent with the requirement to comply with throughput guidelines, and the high connectivity later in the day may correspond to the creation of loans in the overnight interbank market, as well as the start of the American business day. The relatively low connectivity early in the day needs to be interpreted with care: while the level of activity may appear to be lower than later in the day, the payments being made are likely to be 'time-critical', in particular pay-ins to CLS. The apparently slow build-up may result from a lack of payment instructions at this time, but may also reflect a tendency to delay other payments until after time-critical payments have been made. So the impact of an operational disruption at a given node at this time of the day may be no less significant than at times when the network is busier.

6 Impact of an operational event on network topology

The study of network topologies can contribute to an understanding of the stability and robustness of a network of payment flows in response to an operational disturbance. Different network properties may give rise to differing degrees of resilience to such disturbances. Concepts developed in the fields of statistical mechanics and social network analysis – such as those in Newman (2003), for example – can help us to analyse the robustness of payment flows. For example, Albert *et al* (1999), and Crucitti *et al* (2004) find that scale-free networks are robust to random failures, but vulnerable to targeted attacks.(14)

In particular, the properties of a payment network may have important implications for the flow of liquidity through the system in stressed circumstances, for example when a bank is operationally unable to make payments. The higher the connectivity of the system, the faster we might expect liquidity to flow to the stricken member, potentially creating a liquidity sink.(15) Banks that exhibit a low in degree are likely to be more vulnerable to disturbances than other banks, as the removal of one link will severely limit the flow of incoming funds. Conversely,

---

(13) Links with CLS are only active during the settlement window between 06:00 and 11:00.

(14) A scale-free network is a network in which the degree distribution follows a power law relationship; the implication of this is that the network will have the same properties irrespective of the number of nodes. See Soramäki *et al* (2007).

(15) A 'liquidity sink' is the position in which all of a payment system's liquidity ends up in one bank, which is unable or unwilling to recirculate it by making payments itself. This has the result that eventually no other bank will be able to make payments for want of liquidity. The probability of a liquidity sink developing will also be a function of institutional factors: for example, the CHAPS guidelines state that if a member notifies CHAPSCo of a significant problem affecting their capacity to send payments, other members are required to withhold payments to avoid the creation of a sink. Over our sample, this happened 79 times; though almost all of these were for short periods of time only, unlike the operational event we consider.

banks with high out degrees have, *ceteris paribus*, the potential to affect more counterparties if their payment processing is disrupted. However, as we will discuss below, link *weights*, rather than node degree and connectivity, play a larger role in a near-complete network.

In order to get a feel for the robustness of the CHAPS Sterling network, we conduct a case study. In particular, we consider the impact of an operational outage affecting a single bank – hereafter the 'stricken bank' – in 2005. On this day, the stricken bank was able to receive but unable to send CHAPS payments from the start of the day until the middle of the afternoon; that is, until the middle of the afternoon, the stricken bank accumulated liquidity. To put this disruption into some more context, the bank, despite having a large node weight, was actually a net receiver of liquidity on the day in question and the problem was resolved by the end of the day. Both these facts meant that by the end of the day, as we show in what follows, all payments were made and there was no overall liquidity shortage. The impact could have potentially been worse if the bank had been a net supplier of liquidity and if the problem had not been resolved on that day, although CHAPS Sterling has robust procedures in place for dealing with such eventualities. For further details, see Bedford *et al* (2005).

In what follows, we are particularly interested in the effects of this outage on the non-stricken banks. In the cast of a robust network, we would expect such an outage to have little effect on the other banks; that is, we would expect the network of non-stricken banks to have similar properties on the day of the outage as on a normal day.

6.1 Impact on payment activity

An operational disruption impacts most obviously on the intraday profile of payments within the network. Charts 11 and 12 show the average link values and volumes, respectively, for all banks on day *t*, the day of the operational incident. Charts 13 and 14 show the peak values for individual links. (The solid lines indicate the average for the month during which the operational incident occurred, excluding the previous working day, *t*-1, day *t* itself and the following working day, *t*+1, and two standard error bands around this. As there were no other extensions of the system's opening hours during the month, this line drops to zero after 16:30.) The charts indicate that payments were shunted towards the end of the day; once the stricken bank had resolved the problem, payment volumes and values rose as a whole day's worth of payments were made to, and received from, other banks. Indeed, the opening hours of the system were extended to accommodate this.

Chart 11: Average link value for all banks on day *t* and average for month

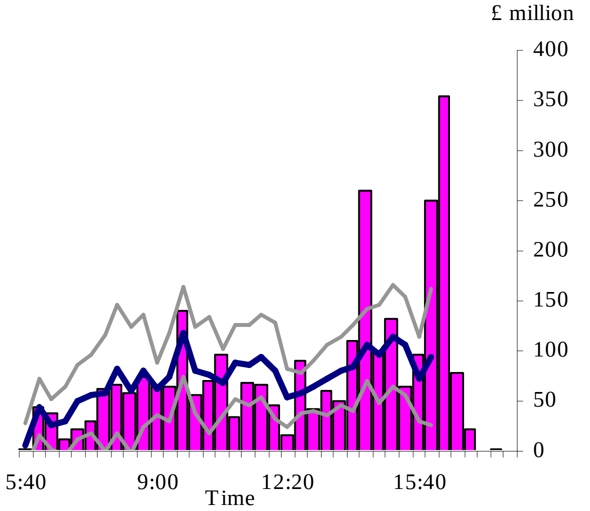

Chart 12: Average link volume for all banks on day *t* and average for month

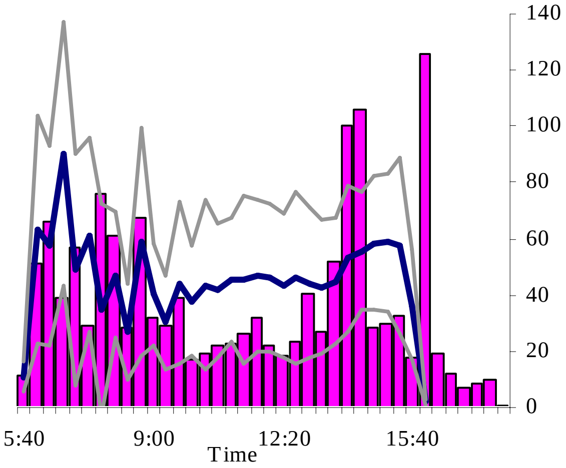

Chart 13: Peak link value for all banks on day *t* and peak for month

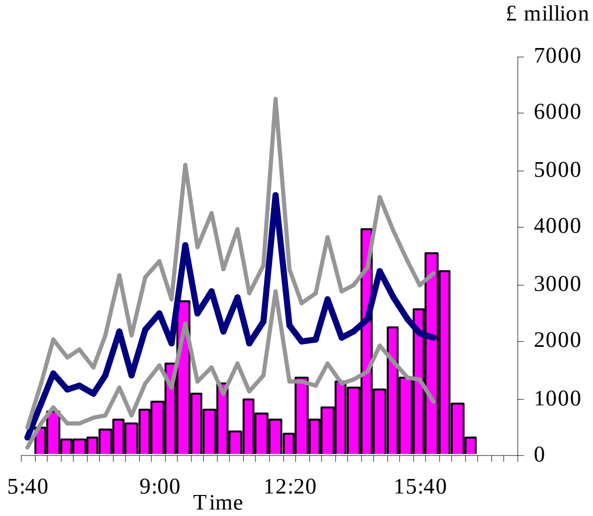

Chart 14: Peak link volume for all banks on day *t* and peak for month

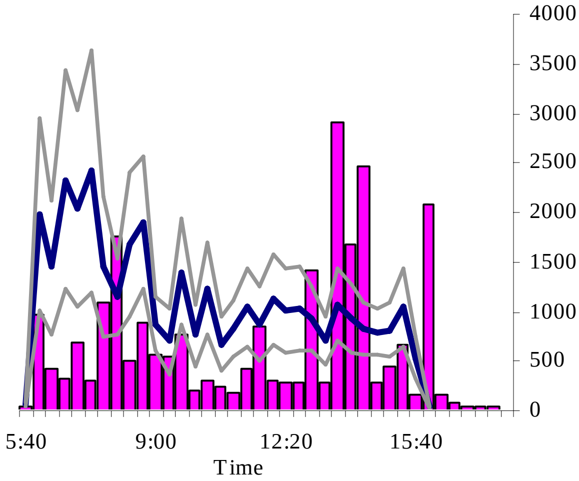

We also see that the disruption on day *t* spilt over to affect payment activity on the following day. On day *t*+1, all the banks held back making large-value payments to the previously stricken bank until 09:00, presumably so they could be sure there would be no repeat of the problem. This is illustrated in Charts 15 to 18. (Again, the solid lines give the average link values and volumes for the month during which the operational incident occurred, excluding the previous working day, *t*-1, day *t* itself and the following working day, *t*+1, and two standard error bands around this.)

Chart 15: Average link value to the day *t* stricken bank on day *t*+1 and average for month

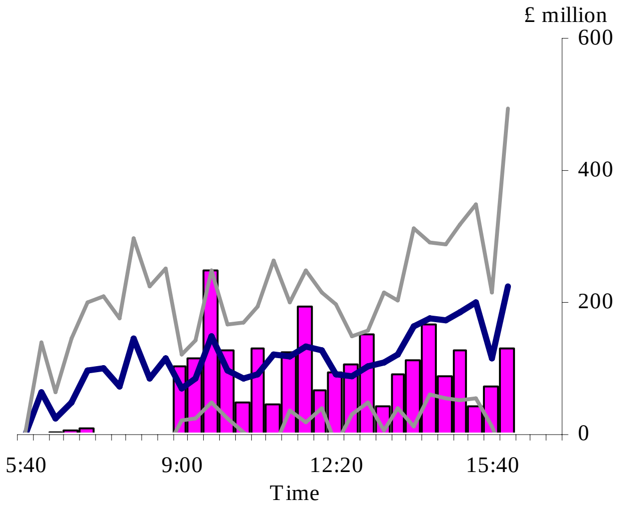

Chart 16: Average link volume to the day *t* stricken bank on day *t*+1 and average for month

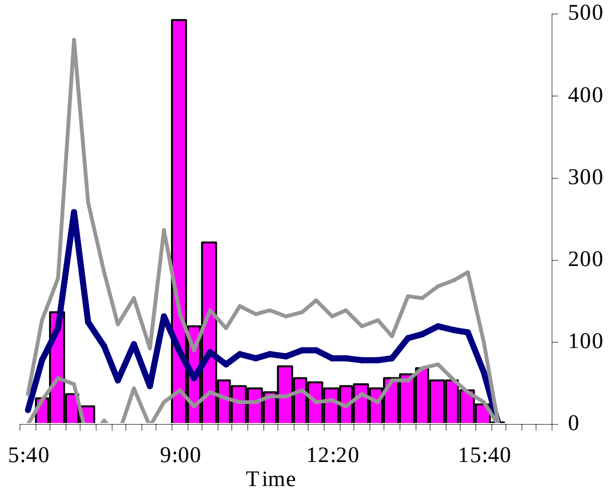

Chart 17: Peak link value to the day *t* stricken bank on day *t*+1 and peak for month

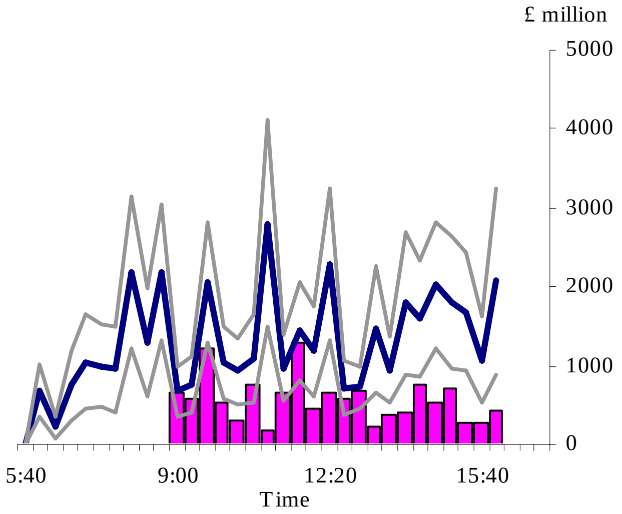

Chart 18: Peak link volume to the day *t* stricken bank on day *t*+1 and peak for month

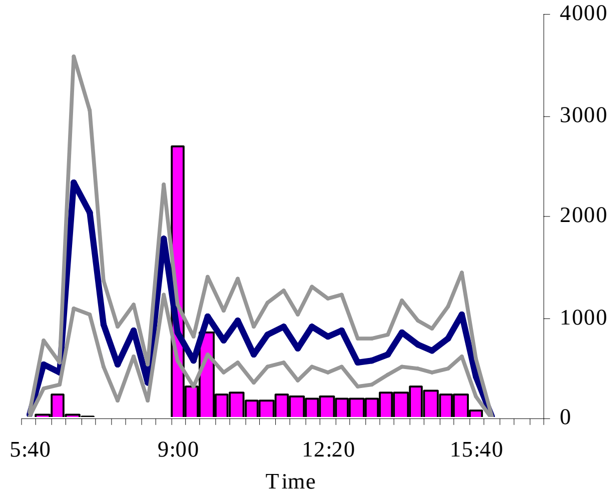

To what extent did this disruption affect the properties of the network of CHAPS payments? Table C contains network statistics for those three dates, together with the averages (and standard deviations in brackets) over our sample as a comparator. The descriptive statistics indicate that the operational outage did not have a significant impact on the properties of the network when viewed at the end of the day. Volumes and values processed fell within the normal range, and the level of connectivity, the average out degree and the average path length remained similar to those of a typical day. But, as the change in the intraday profiles suggests, the end-of-day statistics may conceal important variations in the intraday development of the network. The impact of the outage is therefore more likely to be seen in the properties of the payment network at particular times of the day, both because of the atypical behaviour of the stricken bank, and of the reaction of other banks to this behaviour.

Table C: Network properties around the date of the operational outage

		t-1	*t*	*t*+1	Average
Value out (£ billion)	Total	202.1	204.4	168.8	207.9 (28.2)
	Mean	14.4	14.6	12.1	15.8 (1.9)
Volume out	Total	170,093	104,371	93,431	123,277 (30,751)
	Mean	12,150	7,455	6,674	9,414 (2,356)
Nodes		14	14	14	13 (1)
Links		156	160	158	148 (14)
Average out degree		11.1	11.4	11.3	11.2 (0.2)
Average path length		1.14	1.12	1.13	1.07 (0.06)
Average clustering coefficient		0.90	0.92	0.91	0.95 (0.04)
Connectivity		0.86	0.88	0.87	0.83 (0.06)

Note: Standard deviations in brackets.

To examine this, we calculate network statistics for 'snapshots' of the network at 20-minute intervals throughout the day, using the approach described in Section 5. Charts 19 to 21 display the network statistics for day *t* alongside the average network properties over that particular month (including the stricken bank). The charts also show the minimum and maximum average values over all banks during the respective intraday periods on the whole month (excluding day *t*). As can be seen most of the variation around the mean during day *t* did not exceed these bounds.

Chart 19: Intraday network snapshots: connectivity

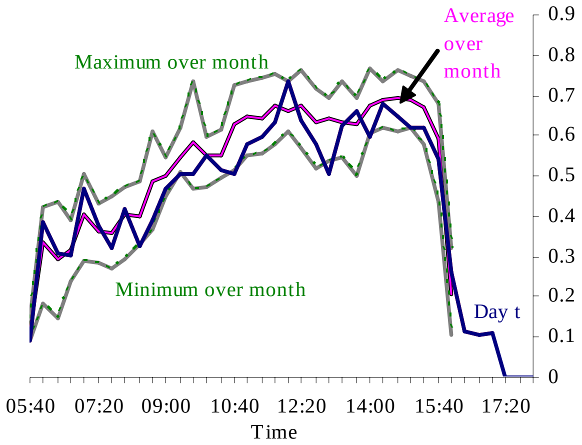

Chart 20: Intraday network snapshots: average path length

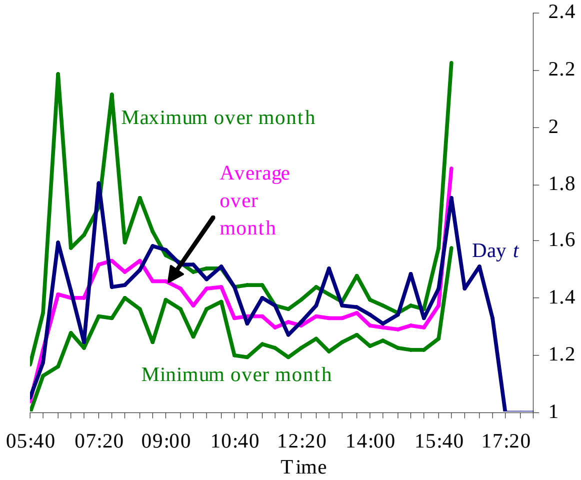

Chart 21: Intraday network snapshots: out degree

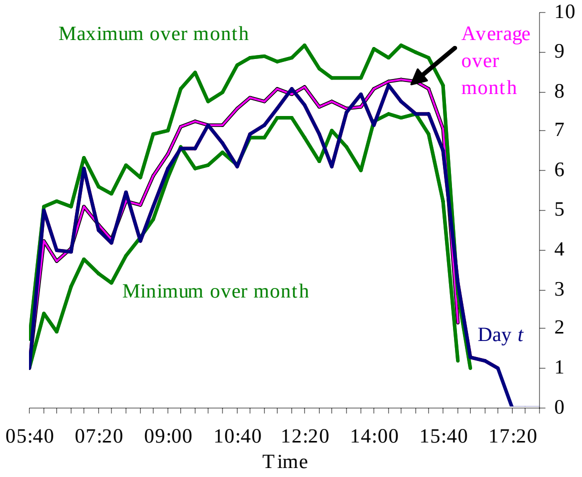

It is notable that the removal of a single node for part of the day had little impact on network properties. Connectivity and average degree were only slightly lower for much of the day than the average for the month. This suggests that the observed impact was down to the fact that the stricken bank was unable to make payments, and that certain other banks ceased to make payments to the stricken bank; there is no obvious sign of disruption in the network of non-stricken banks. Consistent with this, the average path length did not increase appreciably despite the missing node.

To consider directly the impact of the operational incident on payments between the non-stricken banks, Charts 22 and 23 show the average link value and volume, respectively, for the network composed only of these banks. Charts 24 and 25 show the peak values and volumes, respectively, for this same network. (Again, the solid lines give the average link values and volumes for this network for the month during which the operational incident occurred, excluding the previous working day, *t*-1, day *t* itself and the following working day, *t*+1, and two standard error bands around this.) We see that average link values and volumes for payments made between the non-stricken banks were not very different from the average and, by comparing with Charts 11 and 12, that the higher-than-usual values and volumes late in the day were solely attributable to payments to and from the stricken bank.

Chart 22: Average link value for non-stricken banks on day *t*

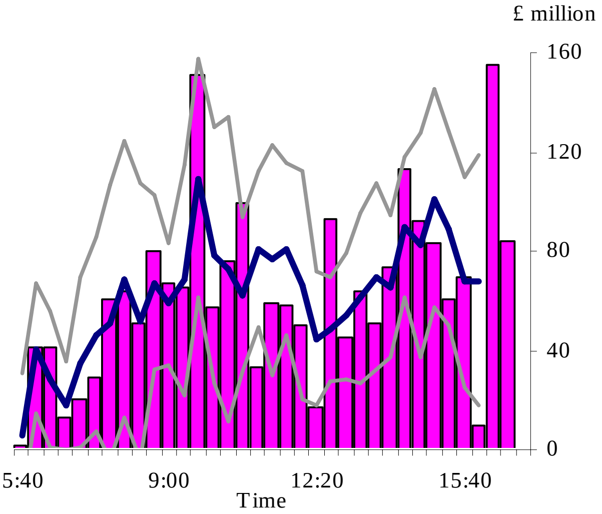

Chart 23: Average link volume for non-stricken banks on day *t*

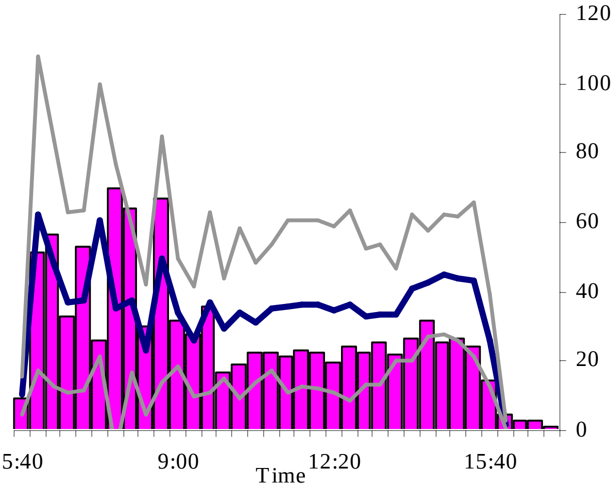

Chart 24: Peak link value for non-stricken banks on day *t*

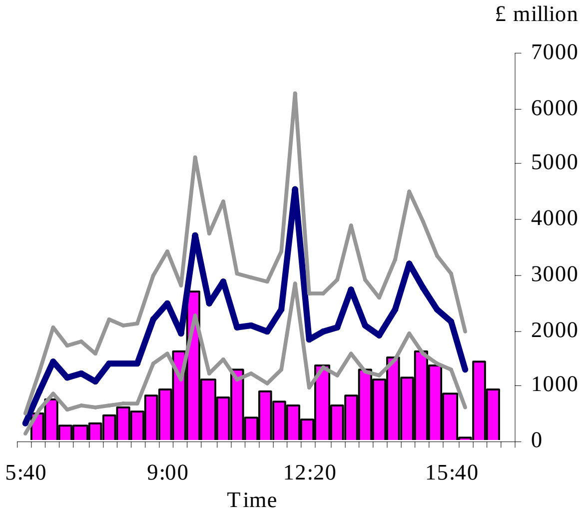

Chart 25: Peak link volume for non-stricken banks on day *t*

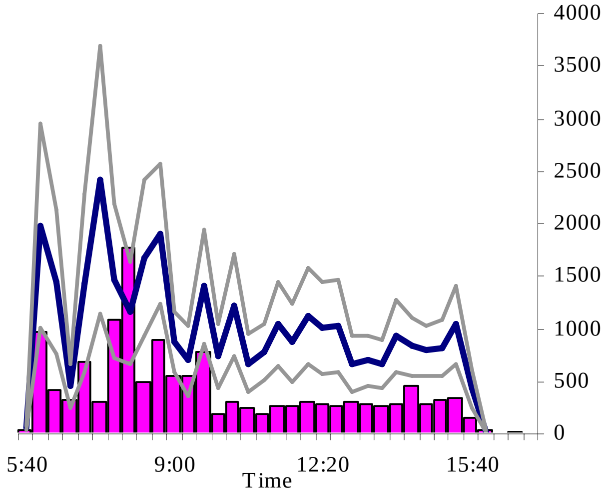

We can further explore these observations by considering a ‘graph edit distance’ measure, which seeks to assess the ‘similarity’ of payment flows around the period of the event to normal payment flows. Such measures generally associate a cost function for each operation required to transform a graph into another, ie how much do you have to do to one set of payments to make them look similar to another set of payments. The minimum edit cost is a measure of difference, or change, between two graphs (see eg, Bunke (1997)). In this case, we use a simple cost function in which the cost equals the changes that are required in link weights to convert the graph to an average graph over the observation period, ie we ask by how much we have to change the value of payments made between each pair of banks around the time of the event operational incident to make them the same as on that time on an average day.

More specifically the edit distance on a given day is calculated as:

$$\varepsilon = \sum_i \sum_j \left| w_{i,j} - \overline{w}_{i,j} \right| \tag{1}$$

where $w_{i,j}$ is the share of bank *i*'s payments to bank *j*, and $\overline{w}_{i,j}$ its average over the observation period. We calculate the measure for three time periods: for the whole day, for payments exchanged during the disruption, and for those exchanged after the stricken bank had resumed payments. We look at these both for the networks of all banks and for networks consisting only of non-stricken banks, and compare them with results for the same time periods on other days: in particular, we consider the month during which the disturbance occurred. The results are shown in Charts 26-28. The measure incorporates information on relative shifts in payments flows and ranges from zero to unity (where a value of zero reflects isomorphism of the two graphs).

Chart 26: Edit distances over whole day

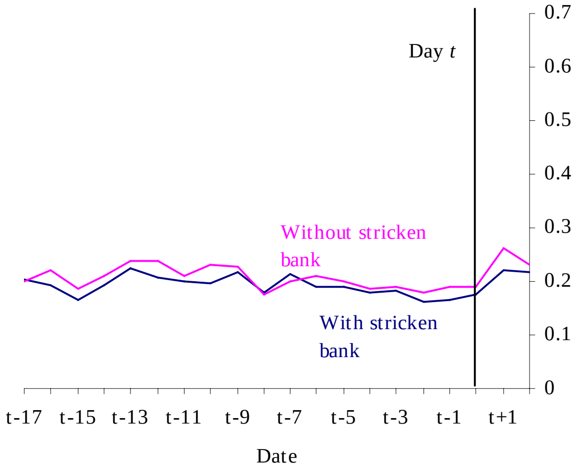

Chart 27: Edit distances during disruption period

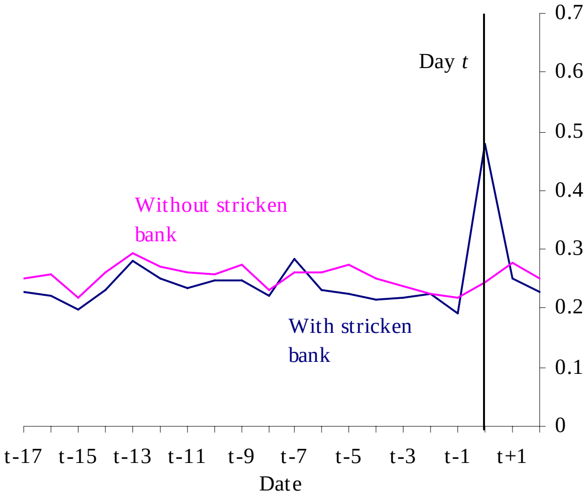

Chart 28: Edit distances after disruption period

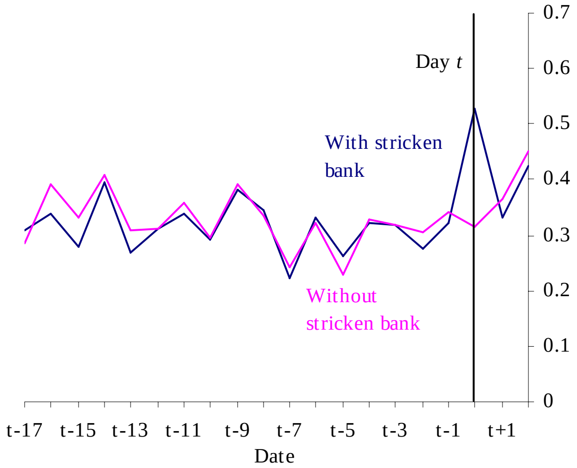

The results of this analysis reinforce the conclusions derived from the earlier analysis. While the whole network is clearly affected by the operational incident, this can be attributed primarily to the impact on the individual bank experiencing the disruption. Payment flows in the network *including* the stricken bank are appreciably different during and after the disruption, but appear normal when viewed over the whole day. And the impact of the outage on the network of non-stricken banks falls within the 'typical' range.

6.2 Impact on system liquidity

The impact of an operational outage at a single node on a network is not merely a matter of the technical ability of other banks to make payments to one another. To the extent that other banks were reliant on the liquidity provided by the stricken node, the removal of this node from the network may have affected their ability to make payments by depriving them of the liquidity needed to make those payments. In other words, even if the outage had little impact on the connectivity of the network, the ability of non-stricken banks to make payments may have been curtailed by the stricken node acting as a liquidity sink. This will depend in part on the reactions of other banks to the outage, that is, on whether and for how long they continued to send payments to the stricken bank.

We can examine directly the response of other banks to the realisation that the stricken bank had stopped sending payments by examining their 'link deficit' towards the stricken bank, that is, the bilateral net sending position with the stricken bank. This is shown in Chart 29, where a positive value indicates that the bank has sent a greater value of payments to the stricken bank than it has received from it, and the solid line shows the average bilateral position over the month during which the operational incident took place. We see that until 14:00 the other banks (taken as a whole) kept sending payments to the stricken bank until the net deficit (roughly) reached the average for the month; this suggests the possibility that the banks applied bilateral net sending limits and kept sending payments to the stricken bank until these limits were reached.(16) After 14:00, the average deficit increased to much more than the monthly average, suggesting that banks began to send payments that they had previously held back as the stricken bank recovered its capacity to make payments.

We can analyse further the liquidity impact of the outage by examining the *multilateral* node deficit for all the non-stricken banks, that is, the total liquidity used by all the non-stricken banks. Chart 30 shows that, for the majority of the day, liquidity usage was no higher than on a 'typical' day. In other words, the fact that one bank was unable to send payments did not mean that the other banks had difficulty sending payments to each other. As evidenced by Chart 23, in which we found that link deficits with respect to the stricken bank did not greatly exceed their average levels for much of the day, this may have resulted from the use of bilateral net sender limits, ie, members were able to actively monitor the payment behaviour of the stricken bank and to withhold payments once it ceased making payments. The evidence in Chart 30 is particularly important since, for the operational incident to have a significant impact, it would need to result in a general liquidity shortage; ie, the amount of liquidity used by the

---

(16) It would be instructive to examine link deficits bank by bank to see if there is evidence that all banks applied bilateral sending limits or that some did and some did not. We leave this for future work.

non-stricken banks serves as an indicator of the extent to which the operational incident is having a negative impact on the rest of the system. The longer the incident went on, the higher the liquidity we would have expected to see used by the non-stricken banks.

Chart 29: Average link net deficit towards the stricken bank on day *t*

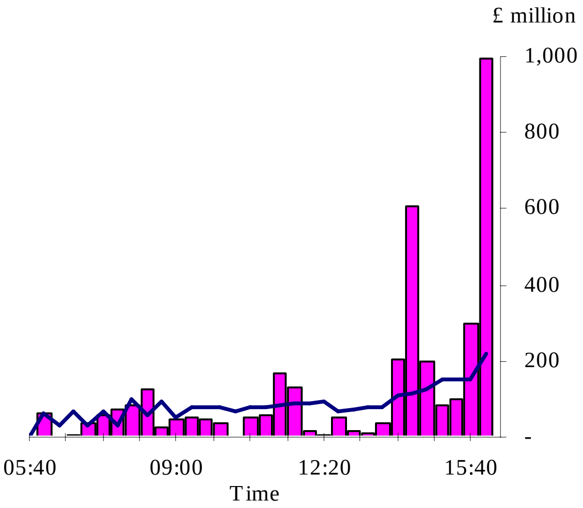

Chart 30: Average link net deficit on day *t* excluding the stricken bank

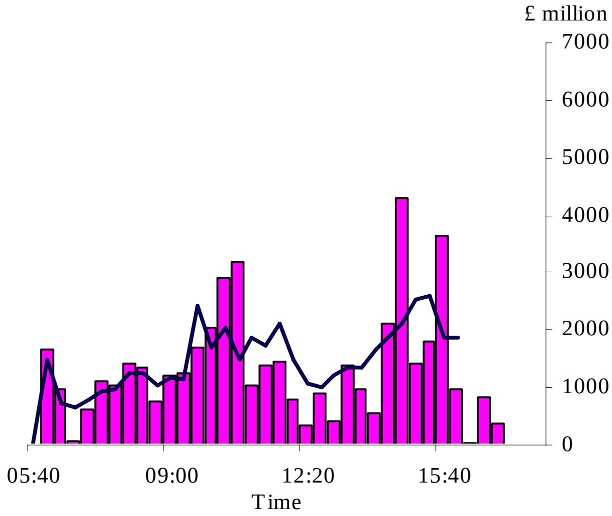

6.3 *Summary*

To summarise, we found the following:

- Since the operational problems were resolved during the day, the incident had no effect at all on the properties of the CHAPS Sterling network viewed from an end-of-day perspective.
- Payments to and from the stricken bank were shifted towards the end of the day.
    - Clearly payments from the bank could not happen until the operational problems were resolved.
    - We found some evidence that the non-stricken banks continued to make payments to the stricken bank until they hit their bilateral sending limits at which point payments were held back until the operational problem was resolved.
- Payments among the non-stricken banks were largely unaffected by the operational incident.
- The non-stricken banks did not obviously require more liquidity in order for this to happen; that is, the stricken bank did not become a ‘liquidity sink’. The use of bilateral sending limits probably helped make sure that the non-stricken banks had enough liquidity to continue making payments to the other non-stricken banks.
- On the following day, the non-stricken banks held back payments to the stricken bank until they were convinced that the operational problems would not repeat.

Taken as a whole, these results broadly indicate that the CHAPS network is robust to this type of shock. This is consistent with the results of the simulation analysis reported in Bedford *et al* (2005). It is also consistent with the general observation that CHAPS is a liquidity-rich system: while incoming funds are frequently used as a funding source, we see that many banks supply sufficient liquidity to insulate themselves from shocks which diminish the flow of incoming funds.(17) Nevertheless, while on this occasion the failure of a single bank to make payments did not have a severe impact on the ability of other banks to make payments between themselves, this may not always be the case. We have shown that there is wide variation in node strength, and as such the impact of an outage may depend critically on which node fails. Moreover, the stricken bank was on this occasion a net *receiver* of liquidity. If the bank in question had been a net *supplier* of liquidity to the system, the impact of the outage on the network may have been more pronounced.

7 Conclusions

The network topology of CHAPS Sterling can be viewed at two levels: the level of the banks that ultimately originate and receive payments or the level of the settlement bank members of the system. Payment flows between banks that originate and receive large-value payments on their own account or on behalf of customers form a complex network, in which payments are routed via one or more of the settlement banks. While considerably smaller than the Fedwire network, some of the characteristics of this network are similar to that created by payments in Fedwire. However, the risks associated with the highly tiered CHAPS structure are different to those of Fedwire, in which all participants are direct members of the system (even though some choose to route payments via a correspondent bank).

This broad network can be collapsed into a smaller network of payment flows between the direct members of the system: the settlement banks. Our analysis of the network topology reveals that payment flows between settlement banks form a near-complete, well-connected network. Most banks make payments directly to all other banks in the system, hence the average path length is close to one and the proportion of potential links actually utilised (the connectivity of the network) is high. We have also seen that the network develops only gradually in the early hours of opening and does not achieve peak connectivity until the early afternoon. The explanation for this pattern lies in the purposes of payments being made at this time, and in particular the tendency to withhold payments until time-critical payments have been settled.[18] We also saw slight peaks in ‘shapshot’ connectivity before noon and late afternoon, indicating that the network is particularly busy at these times. This variation indicates that the impact of an operational disruption may vary according to the time of day at which it strikes.

A consequence of this network structure is that liquidity is able to flow efficiently around the network. The network appears to be highly resilient to an operational shock affecting one of the nodes, corroborating the results of the simulation analysis in Bedford *et al* (2005). In the particular instance of an operational outage examined here, the effective removal of a node for much of the day had little impact on the ability of other banks to make payments between one another. The high connectivity of the network will have contributed to this resilience.

---

(17) See, for example, the discussion in Becher *et al* (2007).

(18) Of course, this tendency to withhold payments means that throughput is less than it otherwise might be. One possible response to this would be to enable banks to ‘ring-fence’ liquidity within the system, ensuring that they always had enough liquidity to meet time-critical payments while ensuring that other payments were still getting made.

As we have discussed, however, the impact of an operational shock may be felt not just on the connectivity of the network but rather on the availability of liquidity with which to make payments. One potentially undesirable feature of a well-connected network is that liquidity may flow more rapidly into a liquidity sink in the event of an operational shock. In the case study, we saw that (some) banks continued to make some payments to the stricken bank, but withheld other payments until the stricken bank regained operational capacity. This is indicative both of the ability of participating banks to monitor each others' behaviour (aided by the small membership of CHAPS) and, perhaps, the use of bilateral net sender limits.

We have also seen that, on this occasion, the failure of the stricken bank to supply liquidity to the system did not disrupt payments elsewhere in the network. This may simply reflect the observation that liquidity is plentiful in CHAPS and as such banks are able to make use of alternative sources of liquidity when the flow of incoming funds is unexpectedly diminished. However, we cannot discount the possibility that the operational disruption of one or more large nodes, especially if they are net suppliers of liquidity to the system, would have a more severe impact on the payment network than was evident from this case study.

References

Albert, R and Barabasi, A L (2002), ‘Statistical mechanics of complex networks’, *Review of Modern Physics*, Vol. 74, pages 47-97.

Albert, R, Barabasi, A L and Jeong, H (1999), ‘Error and attack tolerance of complex networks’, *Nature*, Vol. 406, pages 378–82.

Becher, C, Galbiati, M and Tudela, M (2007), ‘The timing and funding of CHAPS Sterling payments’, *Federal Reserve Bank of New York Economic Policy Review,* forthcoming.

Bedford, P, Millard, S P and Yang, J (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.

Boss, M, Elsinger, H, Summer, M and Thurner, S (2003), ‘The network topology of the interbank market’, *Santa Fe Institute Working Paper 03-10-054.*

Bunke, H (1997), ‘On a relation between graph edit distance and maximum common subgraph’, *Pattern Recognition Letters*, Vol. 18, pages 689-94.

Crucitti, P, Latora, V, Marchiori, M and Rapisarda, A (2004), ‘Error and attack tolerance of complex networks’, *Physica A,* Vol. 340, pages 380-94.

Harrison, S, Lasaosa, A and Tudela, M (2005), ‘Tiering in UK payment systems: credit risk implications’, Bank of England *Financial Stability Review*, December, pages 63-72.

Inaoka, H, Ninomiya, T, Shimizu, T, Takayasu, H and Taniguchi, K (2004), ‘Fractal network derived from banking transaction - an analysis of network structures formed by financial institutions’, *Bank of Japan Working Paper no. 04-E-04.*

James, K (2003), ‘A statistical overview of CHAPS Sterling’, Bank of England *Financial Stability Review*, June, pages 115-21.

Lublóy, Á (2006), ‘Topology of the Hungarian large-value transfer system’, *Magyar Nemzeti Bank Occasional Paper no. 57.*

Newman, M E J (2003), ‘The structure and function of complex networks’, *SIAM Review,* Vol. 45, pages 167-256.

Soramäki, K, Bech, M L, Beyeler, W E, Glass, R J and Arnold, J (2007), ‘The topology of interbank payment flows’, *Physica A*, Vol. 379, pages 317–33.