SAND2007-2930P

New approaches for payment system simulation research

Kimmo Soramäki[1]

Walter Beyeler[2]

Morten Bech[3]

Robert Glass[2]

16 April 2007

Abstract

This article presents new directions for simulation research in interbank payment systems that integrates network topology, network dynamics and agent based modeling of bank behavior. In the process it also reviews literature in the field, and presents applications of the ideas presented. While the focus of the article is on systemic risk in interbank payment systems, the concepts and models presented are applicable to address questions related to other payment systems, and topics such as liquidity flow efficiency as well.

The authors would like to thank Marco Galbiati (Bank of England) and Jeffrey Arnold (Federal Reserve Bank of New York) for comments and their input in the papers summarized in this article.

[1] Helsinki University of Technology, Finland. Corresponding author: e-mail: kimmo@soramaki.net.

[2] Sandia National Laboratories, Albuquerque, NM, USA. Supported through the National Infrastructure Simulation and Analysis Center (NISAC), a program of the Department of Homeland Security's Infrastructure Protection/Risk Management Division and comprised of a core partnership of Sandia National Laboratories Los Alamos National Laboratory. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000."

[3] Federal Reserve Bank of New York, New York, NY, USA. The views expressed in this paper 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.

1 Introduction

At the apex of the financial system is a network of interrelated financial markets by which domestic and international financial institutions allocate capital and manage their exposure to risk. Critical to the smooth functioning of these markets are a number of financial infrastructures that facilitate clearing and settlement. The events of September 11, 2001 underscored both the resiliency and the vulnerabilities of these financial infrastructures to wide-scale disruptions. Any interruption in the normal operations of these infrastructures may seriously impact not only the financial system, but also the economy as a whole.

A growing body of policy oriented research is available. One segment of the literature focuses on simulating the default of a major participant and evaluating the effects on other institutions in payments[1] and securities settlement systems[2]. Another segment presents detailed case studies on the responses of the U.S. financial system to shocks such as the 1987 stock market crash and the attacks of September 11, 2001.[3] Much of the research has been conducted using data from real operating environments with the given payment flows and settlement rules of the respective systems. As such they are useful for assessing the operation of the particular system under disruptions, but the results are difficult to generalize to systems with other characteristics. Little research has focused on explaining the relationship between the characteristics of the system and its performance during and following disruptions. Also the behavior of participants has been generally exogenously defined or assumed unchanged (or to change in a predetermined manner) when the policy parameters of the system change, or when a bank changes its settlement behavior as a consequence of operational of financial problems. Such assumptions are unlikely to hold in the case in real disruptions.

This article argues that three aspects are important for answering the still unanswered questions on what makes a payment system and its participants robust or fragile towards disruptions, and what are the most efficient measures to reduce the likelihood and magnitude of disturbances. First, understanding the pattern of liquidity flows among the system participants. Second, understanding how the rules of the system affect the dynamics of liquidity flows. Third, the ability to evaluate likely behavioral changes of the participants before, during and following disruptions or as a consequence of policy changes.

This article presents new approaches at answering the above questions. It is organized as follows. Section two discusses how payment system interactions can be described by means of network topology and presents empirical results for the US Fedwire system. Section three describes dynamics that can take place in interbank payment systems and presents a simple model of a payment system based on simple rules of settlement. Section

[1] see Humphrey (1986), Angelini et al. (1996), Kuussaari (1996), Bech et al. (2002), Northcott (2002), Bech and Soramäki (2005), Bedford et al. (2005), and Mazars and Woelfel (2005)

[2] see Hellqvist and Koskinen (2005) and Devriese and Mitchell (2006).

[3] see Bernanke (1990), McAndrews and Potter (2002) and Lacker (2004).

four presents some possible directions for modeling participant behavior in payment systems. Section five concludes.

2 Modeling interbank payment flows

A payment system can be treated as a specific example of a complex network (see e.g. Newman 2003). In recent years, the physics community has made significant progress towards understanding the structure and functioning of complex networks. The literature has focused on characterizing the structure of networked systems and how the properties of the observed topologies relate to stability, resiliency and efficiency in case of perturbations and disturbances.

From a technical perspective, most payment systems are star networks where all participants are linked to a central hub (the operator) via a proprietary telecommunications network. From a payment processing perspective, payment systems are generally complete networks as all nodes (participants) are linked in the sense that they can send and receive payments from each other. However, these representations do not necessarily reflect the actual behavior of participants that controls the flow of liquidity in the system and thus the channels for contagious transmission of financial disturbances. In common with other of social networks mediated by technology (such as e-mail or telephone calling), the networks formed by actual participant behavior are of more interest than the network structure of the underlying communication system.

2.1 Network representation of payment systems

Networks have been modeled in several disciplines such as in mathematics and computer science under graph theory, in applied mathematics and physics under network theory and in sociology under social network analysis. While the terminologies and research questions in the different traditions vary, common to all is the representation of the topic under study as (at minimum) two types of elements: nodes and connections between them, i.e., links. The following paragraphs summarize the main concepts.

Links can be either undirected or directed. Links can have weights attached to them representing the importance of the relationship between nodes. The strength of a node can be calculated as the sum of the weights of all the links attached to it. For a directed network, strength can be defined over both the incoming and outgoing links.

A link from a node to itself is called a loop. The neighbors of a node are all the nodes to which it has a link. The predecessors of a node are the nodes that have a link to the node and the successors are the nodes that have a link from the node. A walk is a sequence of nodes in which each node is linked to the next. A walk is a path if all its nodes are distinct. The length of a path is measured by the number links. If the start node and the end node of a path are one and the same, then it forms a cycle.

A complete network is a network where all nodes have a link to each other. A tree is a network in which any two nodes are connected by exactly one path. A connected network is a network where any two nodes can be joined by a path while a disconnected network is made up of two or more connected components or sub-networks. These concepts are illustrated in Fig. 1a.

Figure 1: Network modeling

a) b)

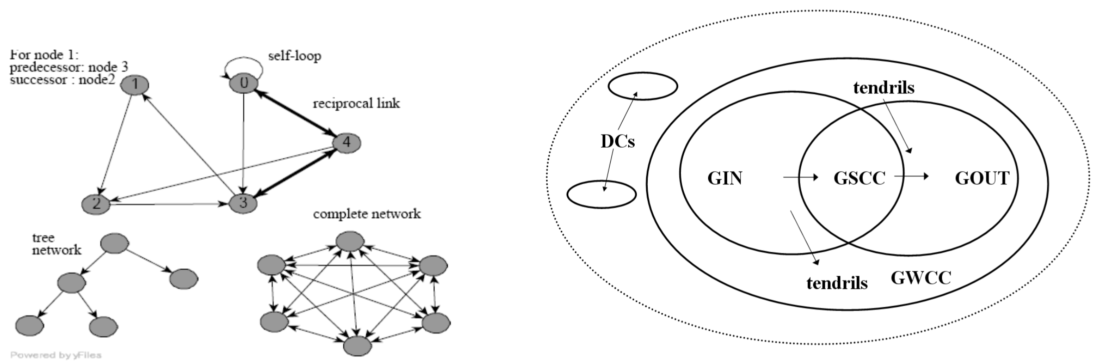

The most basic properties of a network are the number of nodes *n* and the number of links *m*. The number of nodes defines the size of a network while the number of links relative to the number of possible links defines the *connectivity* of a network. The *degree* of the network is the average number of links for each node in the network.

A starting point for the quantitative analysis of a network is to partition the set of nodes into *components* according to how they connect with other nodes. Dorogovtsev et al. (2001) divide a network into a single giant weakly connected component (GWCC) and a set of *disconnected components* (DCs). The GWCC is the largest component of the network in which all nodes connect to each other via undirected paths. The DCs are smaller components for which the same is true. The GWCC consists of a *giant strongly connected component* (GSCC), a *giant out-component* (GOUT), a giant in-component (GIN) and *tendrils*. The GSCC comprises all nodes that can reach each other through a directed path. A node is in the GOUT if it has a path from the GSCC but not to the GSCC. In contrast, a node is in GIN if it has a path to the GSCC but not from it. Tendrils are nodes that have no directed path to or from the GSCC. They have a path to the GOUT or a path from the GIN (see Figure 1b).

Application of the component analyses to liquidity flows between banks provide insights on the structure of these flows within the payment system and give clues with respect to the relative importance and vulnerability of banks in the system in case of disruptions. As banks in GOUT only receive funds from other banks in the GSCC, a disruption by a bank in GOUT would only affect other banks in that component. Banks in GIN are affected only by disruptions in the same component, and not by banks in other components as their payment processing is not dependent on incoming liquidity from these banks. Banks outside the GSCC tend to be smaller whereas all money center banks belong to the GSCC.

Two important characteristics of a node in a directed network are the number of links that originate from the node and the number of links that terminate at the node. These two quantities are referred to as the *out-degree* and *in-degree* of a node respectively. The average degree of a node in a network is the number of links divided by the number of nodes, i.e. $<k>=m/n$. Networks are often categorized by their degree distributions. The degree distribution of a classical random network (ER-network, Erdõs and Rényi 1959) is a Poisson distribution. Many real networks have fat-tailed degree distributions and a large number have been found to follow the power law $P(k_i = x) \sim k^{-y}$ for large-degree nodes. Networks with a power-law distribution are sometimes referred to as scale-free networks[4]. Scale free networks have been found to remain better connected when subjected to random failures than other types of networks. Albert et al. (1999) and Crucitti et al. (2004) find that the connectedness of scale-free networks is robust to random failures, but vulnerable to targeted attacks. However, one must be a bit careful here as the process acting on the network influences such analyses of robustness and vulnerability.

Simply put, banks that have a low in-degree and high weights for these links are likely to be more vulnerable to disturbances than other banks as the removal of one link will severely limit the amount of incoming funds. Conversely, banks with high out degree have ceteris paribus the potential to affect more counterparties if their payment processing is disrupted. Understanding the topology of payment flows is likely to be important in assessing the resiliency of a payment system to wide-scale disruptions.

It is also common to analyze *distances* between nodes in the network. The distance from node i to node j is the length of the shortest path between the two nodes. The average distance from a node to any other node in a strongly connected network is commonly referred to as the average path length of a node. If the network is not strongly connected, paths between all nodes may not exist. In a payment network the path length may be important due to the fact that the shorter the distances between banks in the network, the easier liquidity can re-circulate among the banks. On the other hand, a payment system where liquidity flows over short paths is also likely to be more vulnerable to disruptions in these flows.

Sociologists have long studied *clustering* in social networks, i.e., the probability that two nodes which are the neighbors of the same node, themselves share a link. This is equivalent to the observation that two people, each of whom is your friend, are likely to be friends with each other. One way of measuring the tendency to cluster is the ratio of the actual number of links between the neighbors of a node over the number of potential links among them. A tree network has a clustering coefficient of zero, and a complete network a coefficient of one. In a classical random network, the clustering coefficient is the unconditional probability of connection, i.e. $\langle C \rangle = p$.

In a payment network, the clustering coefficient measures the prevalence of payments between a bank's counterparties. In terms of resilience one could hypothesize that disturbances in banks with a higher clustering coefficient might have a compounding impact on their counterparties, as some of the disturbance may be passed on by the bank's neighbors to each other - in addition to the direct contagion from the source of the disruption.

[4] This is because the power law distribution is the only scale-free distribution, i.e. if the scale by which x is measured is increased by a factor, the shape of the distribution p(x) is unchanged, except for an overall multiplicative constant (see Newman 2005)

There are various measures of the centrality that indicate the relative importance of nodes in a network. Four measures of centrality are commonly used in network analysis: degree, closeness, betweenness, and eigenvector centrality. The first three were described in their current form by Freeman (1979) while the last was proposed by Bonacich (1972). Degree centrality takes into account only the immediate neighborhood of the node, i.e. it is simply the number of links the node has. Closeness centrality as defined by Freeman is the sum of shortest paths from all other nodes. Betweenness centrality may be defined loosely as the number of times that a node is on the shortest path between any pair of nodes. Eigenvector centrality encapsulates the idea that the centrality of a node depends also on the centrality of the nodes that it is linked by (or links to). A famous commercialization of this centrality measure is the PageRank algorithm by Google (Brin and Page 1995). In general, the importance of the node will depend on process taking place in the network. Borgatti (2005) provides a good overview of alternative processes in networks and centrality measures applicable for their analysis.

Finally, a key question in the study of networks is how the topologies that are seen in reality have come into being. There are two classes of network formation models some times referred to as equilibrium and non-equilibrium models (Dorogovtsev and Mendes 2003). Equilibrium models have a fixed set of nodes with randomly chosen pairs of nodes connected by links. Erdõs and Rényi (1959) develop a basic model of a n node network, with each pair of nodes connected by a link with probability p. This type of network is commonly referred to as a classical random network. Non-equilibrium network models grow a network by successively adding nodes and setting probabilities for links forming between the new nodes and existing nodes and between already existing nodes. Many of these models, notably the Barabasi and Albert (1999) model (BA-model), are based on preferential attachment. Preferential attachment assigns a probability of a link forming with a node that is increasing with the number of prior links of the node.

2.2 Fedwire as an example of a complex network

Soramäki et al. (2007) analyze the topology of daily networks formed by the payment flows between commercial banks over Fedwire for a period of 62 consecutive business days. Apart from a few holidays, the statistics characterizing the network were quite similar from day to day. These networks shared many characteristics with other empirical complex networks, such as a scale-free degree distribution, high clustering coefficient and the small world phenomenon (short path lengths in spite of low connectivity). Like many other technological networks, high-degree nodes tend to connect to low-degree nodes. Similar conclusions can also be reached from analysis on BoJ-NET by Inaoka et al. (2005).

Moreover, Soramäki et al. (2007) report that the topology of the network was significantly altered by the attacks of September 11th, 2001. The number of nodes and links in the network and its connectivity was reduced, while the average path length between nodes was significantly increased. Interestingly, these alterations were of both similar magnitude and direction to those that occurred on several of the holidays contained within the period.

Figure 2a shows liquidity flows in Fedwire as a visual graph. The figure includes over 6,600 nodes and more than 70,000 links. Each link between two banks is shaded by the value of payments exchanged between them, with darker shades indicating higher values. Despite the appearance of a giant fur ball, the graph suggests the existence of a small group of banks connected by high value links. To gain a clearer picture of this group, a subset of the network where the focus is on high value links is displayed in Figure 2b. This graph shows the largest undirected links that comprise 75 percent of the value transferred. The network consists of only 66 nodes and 181 links. The prominent feature is a densely connected sub-graph, or clique, of 25 nodes to which the remaining nodes connect. By itself it is almost a complete graph. A small number of banks and the links between them thus dominate the value of all payments sent over the network.

Figure 2: Visualization of the liquidity flow network (Soramäki et al. 2006)

a) b)

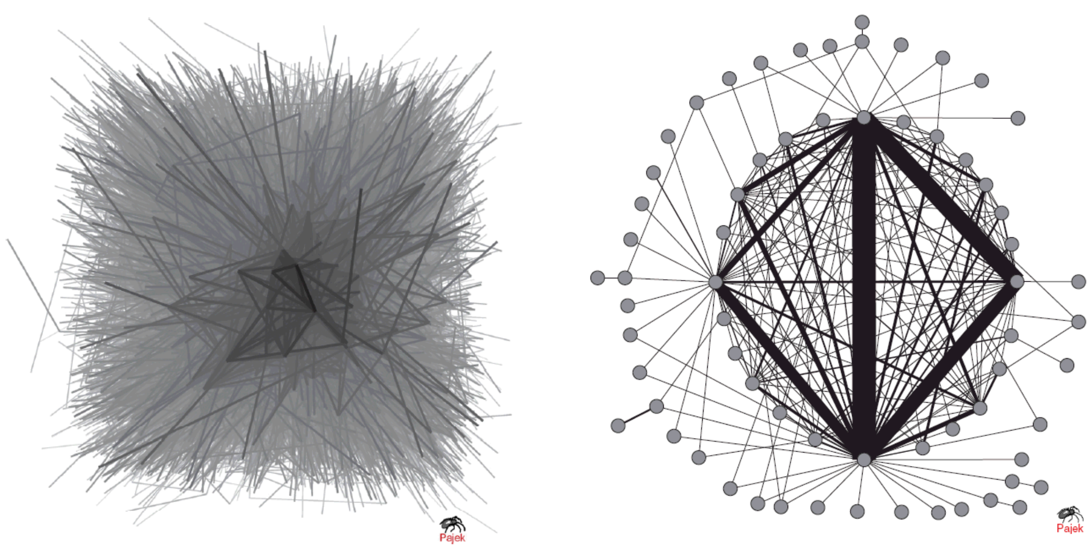

The analysis finds that payment networks have characteristics similar to other social and technological networks. An unanswered question why the network has the structure it does: the network may grow over time by a logic that is very general or that is particular to payment systems, or to specific policies of a given system. This is an interesting topic for future research. The network structure has also implications for its robustness. Robustness of the network, however, also depends on the processes taking place in it. This is the topic of the next sections.

3 Modeling payment system dynamics

3.1 Network dynamics

A number of payment system simulations carried out in recent years have used actual or generated payment data. These simulations have studied the actual dynamics of payment systems, where system rules have varied from simple real-time gross settlement to complex hybrid settlement mechanisms with offsetting and multilateral settlement capabilities. The research can be summarized as trade-off questions between liquidity, speed of settlement and risks. The impact of bank behavior has not been taken endogenously into account in these simulations. A summary of this line of research is provided in Leinonen (2005) and is not presented here.

From a network perspective, the performance of banks (nodes) is often dynamically dependent on the performance of other banks within the network and upon the structure of linkages between banks. A failure by one node in the network, for example, may hinder flows in the network and adversely impact the performance of the other nodes as the disturbance propagates in the network.

One branch of network literature has investigated the resilience of different network topologies in terms of a connectivity threshold (i.e. percolation threshold)[5] at which a network dissolves into several disconnected components. A well known finding is that scale free networks are more robust to random failures than other types of networks. However, they are very susceptible to the removal of the very few highly connected nodes. These static failure analyses may be applicable to some networks if the interest is the availability of paths between nodes in the network - but are less applicable to networks of monetary flows which contain both flows via the shortest paths as well as longer walks within the network.

Another branch of the literature has studied the impact of perturbations that cascade through the network on the basis of established theoretical or domain-specific rules[6]. In these dynamical models nodes generally have a capacity to operate at a certain load, and once the threshold is exceeded, some or all of the node's load is distributed to neighboring nodes in the network (Bak et al. 1987). While the detailed dynamics depend on the rules applied for the cascades, generally the most connected nodes (or nodes with highest load in relation to overall capacity) are more likely than average nodes to trigger cascades. Increased heterogeneity makes the system more robust to random failures, but more susceptible to targeted attacks that may cause global cascades.

Cascade models have been applied by physicists to systems within fields ranging from geology to biology to sociology (e.g., Jensen 1998). This research has demonstrated that models made of very simple agents, interacting with neighboring agents, can yield surprising insights about system-level behavior. In the spirit of these cascade models,

---

[5] e.g. Bollobas (1985), Moore and Newman (2000) and Callaway et al. (2000)

[6] e.g. Watts (2002) and Crucitti et al. (2004b) for random and complex networks, respectively, and Sachtjen et al. (2000) and Kinney et al. (2004) for power networks.

Beyeler et al. (2006) formulate a simple agent-based model for liquidity flows within a payment system.

3.2 Simple payment system model

The model of Beyeler et al. includes only the essential processes of a payment system and its accompanying liquidity market. A set of banks exchange payments through a single common payment system. All payments occur only along the links of a scale-free network - as was shown to be representative of Fedwire liquidity flows. Banks' customers randomly instruct them to make a unit payment to a neighboring connected bank. Banks are reflexively cooperative: they submit the payment if the balance in their payment system account allows; otherwise they place the instruction on a queue for later settlement.

If the receiving bank has instructions in its queue, the payment it just received enables it to remove a queued instruction and submit a payment in turn. If the bank that receives that payment is also queuing instructions, then it can make a payment, and so on. In this way a single initial payment made by a bank can cause many payments to be released from the queues of the downstream receiving banks. This is an example of the cascade processes typically studied in other models of self-organized criticality. Statistics on these settlement cascades are an indicator of the extent of interdependence of the banks, and in the model, they are a controlled by two parameters, the overall liquidity and market conductance.

Figure 3: Simple payment system model (Beyeler et al. 2006)

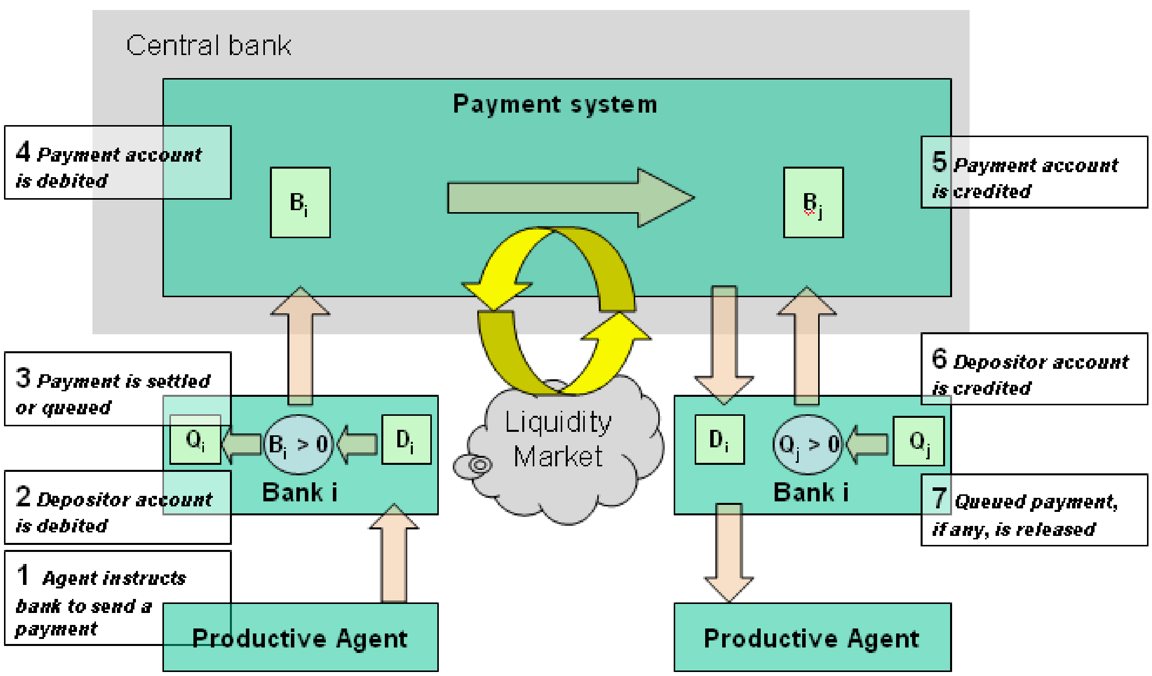

In the absence of a liquidity market, only abundant liquidity allows banks to operate independently; reducing liquidity increases the likelihood that a given bank will exhaust its balance and begin queuing payments. A bank that has exhausted its balance must wait for an incoming payment from one of its neighbors. When liquidity is low a bank's ability to process payments becomes coupled to its neighbors' ability to process. The output of the payment system as a whole is no longer determined by overall input, but instead becomes dominated by the internal dynamics of the system. Figure 4a shows how the correlation between arriving instructions and submitted payments degrades in the model as liquidity is reduced (green: high liquidity; orange: medium liquidity; red: low liquidity). A settlement cascade, that is the release of queued payments as a result of a single initiating payment, can comprise hundreds of queued payments as illustrated in Figure 4b.

To explore how liquidity markets reduce coupling among network neighbors and thereby reduce congestion, market transactions were represented as a diffusive process where a bank's balance plays the role of a potential energy or pressure. Banks with high balances tend to contribute liquidity to the market, while banks with low balances tend to draw liquidity from the market. There is no decision-making or price setting in this simple market model, but it reflects two essential features of a real market: liquidity flows from banks with surplus funds to banks that need funds, and liquidity can flow from any bank to any bank – flows are not confined to the links of the payment network. It creates a separate global pathway for liquidity flow. The ease of liquidity flow through the market is described by a single conductance parameter.

Figure 4: Instruction and Payment Correlation (a), and Settlement Cascade Length Distribution (b).

a)

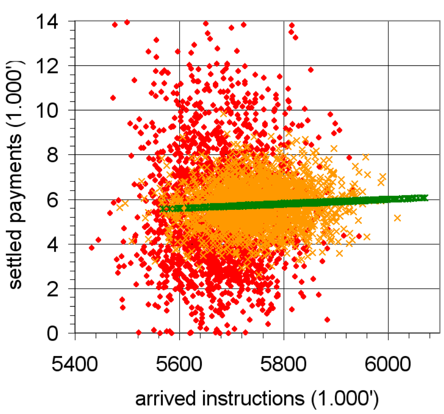

b)

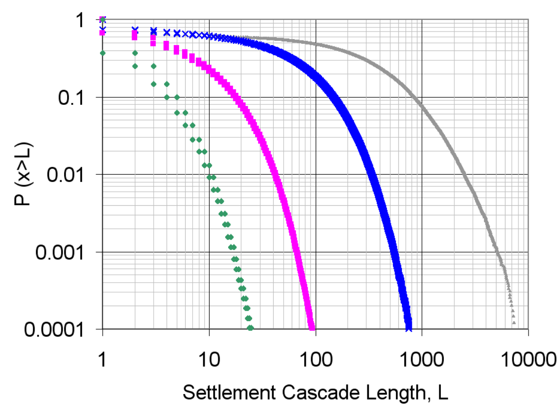

With a liquidity market included, the number of payments closely tracks the number of instructions as the coupling between banks is weakened and the size of the settlement cascades is reduced. The rate of liquidity flow through the market relative to the rate of flow through the payment system was surprisingly small. The performance of the system can be greatly improved even though less than 2 per cent of the system through-put flows through the market.

4 Modeling bank behavior

4.1 Decision making, learning and adaptation

Wide-scale disruptions may not only present operational challenges for participants in the interbank payment system, but they may also induce participants to change the way they conduct business. The actions of participants have the potential to either mitigate or exacerbate adverse effects. Hence, understanding how participants interact and react when faced with operational adversity will assist operators and regulators in designing countermeasures, devising policy, and providing emergency assistance, if necessary.

The first approach to study bank behavior in payment systems has been to use standard game theory. Angelini (1998) and Kobayakawa (1997) use a setup derived from earlier literature on precautionary demand for reserves. Angelini (1998) shows that in a RTGS system, where banks are charged for intraday liquidity, payments will tend to be delayed and that the equilibrium outcome is not socially optimal. Kobayakawa (1997) models the intraday liquidity management process as a game of uncertainty, i.e., a game where nature moves after the players. Kobayakawa (1997) shows that both delaying and not delaying can be equilibrium outcomes when intraday overdrafts are priced. McAndrews and Rajan (2002) study the timing and funding of transfers in the Fedwire funds transfer system. They show that banks benefit from synchronizing their payment pattern over the course of the business day because it reduces the overdrafts. Bech and Garratt (2003) develop a stylized two period - two player model with imperfect information. They analyze the strategic incentives under different intraday credit policy regimes employed by central banks and characterize how the Nash equilibria depend on the underlying cost parameters for liquidity and delays. It turns out that two classical paradigms in game theory emerge: the Prisoner's Dilemma in case where intraday credit is provided against collateral and the Stag Hunt coordination game in the case where the central bank charges a fee. Hence, many policy issues can be understood in terms of well-known conflicts and dilemmas in economics.

Other approaches that have been applied to similar problems of repeated interaction among a large number of players are evolutionary game theory and reinforcement learning (such as Q-Learning by Watkins et al. 1992). Agents who learn about each others' actions through repeated strategic interaction is a leading theme in evolutionary game theory. In most of the existing literature it is customary to look at the players' asymptotic behavior in situations where the payoffs are some known function of players' strategies. In one strand of the literature, this knowledge is a prerogative of the players, who can therefore use adaptive rules of the type "choose a best reply to the current strategy profile". In a second research line, the learning rules do not require knowledge of the payoff function on the part of the learners. Such rules are instead of the kind "adopt more frequently a strategy that has given a high payoff".

Galbiati and Soramäki (2007) use methods from reinforcement learning (Barto and Sutton 1998) and fictitious play (Brown 1951) to numerically solve a model with interactions among a large number of banks that settle payments on a continuous basis under imperfect information, stochastic payoffs and a finite but long sequence of settlement days. The model is summarized and discussed in more detail below.

4.2 Multi-agent model of bank behavior

Galbiati and Soramäki (2007) develop a dynamic multi-agent model of an interbank payment system where payments are settled on the basis of pre-committed funds. In the model banks choose their level of committed funds on the basis of private payoff maximization.

The model consists of a sequence of settlement days. Each of these days is a simultaneous-move game, in which each bank chooses the amount of liquidity to commit for payment processing, and receives a stochastic payoff. Payoffs are determined by means of simulating the settlement day with the amounts of liquidity chosen by the banks. Instructions to be settled by the banks arrive on the basis of a Poisson process and are ex-ante unknown to the banks. As shown in section 3.2, the relationship between instruction arrival and payment settlement is very complex and could not so far be described analytically. Adaptation takes place through reinforcement learning with Bayesian updating, with banks maximizing immediate payoffs. Figure 5 shows the sequence of decisions, events, and learning in the model.

Figure 5: Overview of a multi-agent learning model of a payment system (Galbiati and Soramäki 2007)

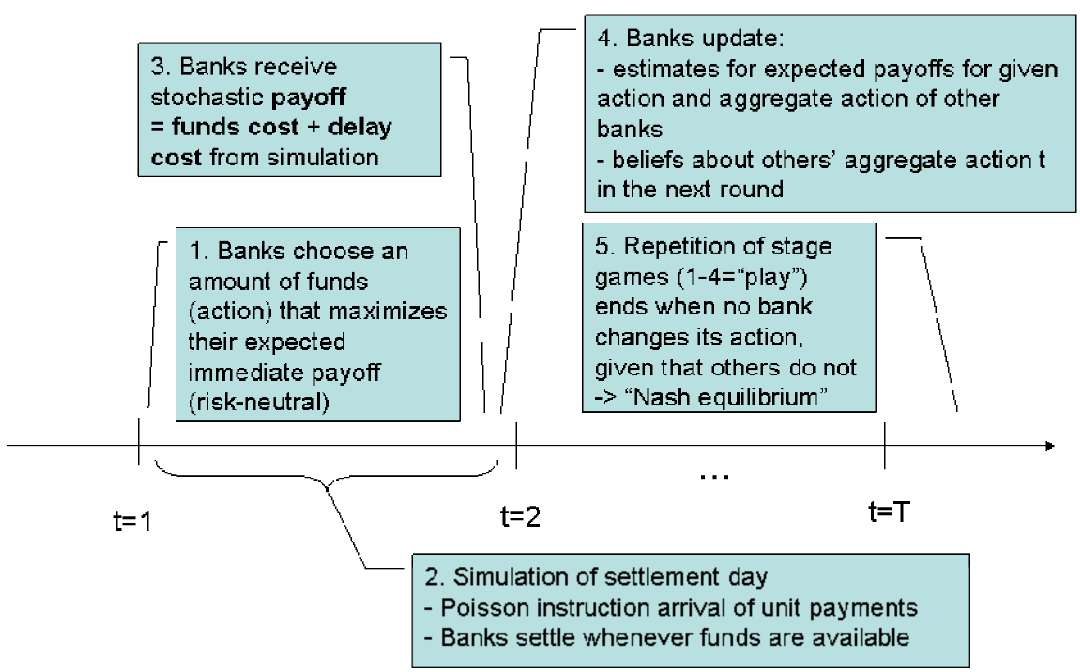

By the process of individual pay-off maximization, banks adjust their demand for liquidity up (reducing delays) when delay costs increase, and down (increasing delays), when they rise. It is well known that the demand for intraday credit is generated by a tradeoff between the costs associated with delaying payments, and liquidity costs. Simulating the model for different parameter values, they find that the demand for intraday credit is an S-shaped function of the cost ratio between intraday credit costs and the costs associated with delaying payments[7] (see Figure 6a).

Figure 6: Demand for intraday credit (a), Payoff comparison (b).

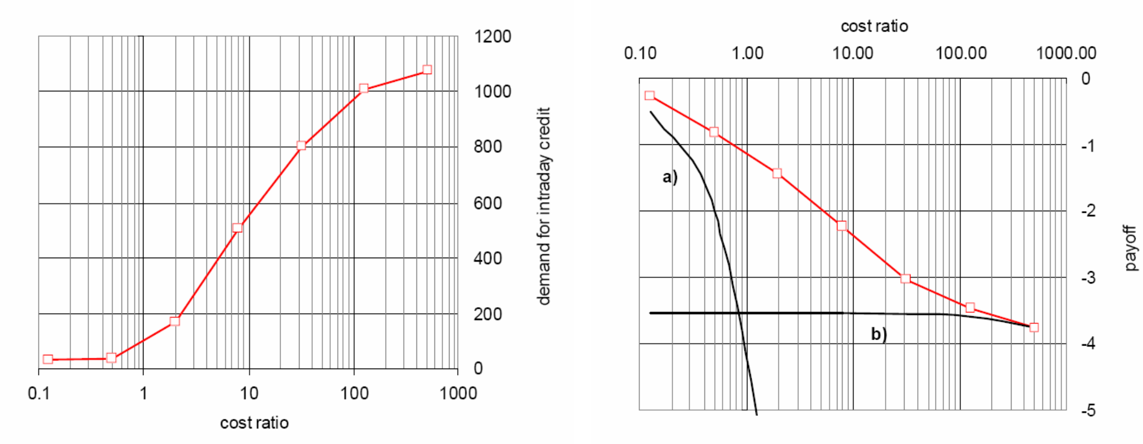

An interesting question is how good the performance of the banks is in absolute terms. To understand this we compare the payoffs received by the banks through adaptation with two extreme strategies:

a) delay all payments to the end of the day;

b) commit enough liquidity to be able to process all payments promptly.

The performance of these three strategies is shown in Figure 6b. For any level of the delay cost, the adaptive banks obtain better payoffs than either of the two extreme strategies, as they manage to learn a convenient trade-off between delay and liquidity costs. On the contrary, the strategy under a) becomes quickly very expensive as delay costs increase, and the strategy under b) is exceedingly expensive when delays are not costly.

Ideally, banks should be taking into consideration the future stream of pay-offs as well. This would create a value of information to the banks as discounting expected future payoffs would create an explicit trade-off between exploitation (the use of actions that appear optimal in the light of the available information), and exploration (the use of seemingly sub-optimal actions, which might appear such because of lack of experimentation). Banks may also be risk-averse, interested not only in the expected pay-off, but also its variability. These are among the topics for future research.

5 Conclusion

This article presented three elements of payment systems, new approaches for understanding and analyzing them, and presented examples on how these approaches can be applied to specific research questions. It argues that performance of a payment system

[7] in the model both costs are assumed to be linear

is a function of network topology, the "physics" of the system and the behavior of banks – one factor alone is not enough to evaluate efficiency or robustness.

First, the payment system can be understood as a network of liquidity flows and can be modeled as a graph. Each model of a payment system assumes some topology, be it random, complete or a topology closer to the system being modeled - such as the scale-free topology of Fedwire. Graph theory and social network analysis provide good tools for analyzing the structure of interbank payment systems and their liquidity flows. Understanding how banks are connected in the payment network is important for analyzing their robustness. The concepts developed in the field can help us structurally analyze payment flows in the system (see e.g. Newman 2003). Measures of average path length can tell us how quickly disturbances are likely to reach other banks in the network. More research is clearly needed to identify measures that explain the connection between system topology and its robustness. Centrality measures can help us identify banks that are not only important through their size, but also due to their position in the network and due to their linkages to other banks (see e.g. Borgatti 2005). A likely fruitful area in payment system research would be to use such approaches for the identification of important (and vulnerable) banks in networks representing RTGS or netting systems.

Second, payment systems have rules, procedures and technical constraints for the processing of individual payments that may produce emergent behavior at the system level. An example of these is the settlement cascades that take place at low levels of liquidity and low market conductance. The model of payment system dynamics exhibits a transition from independent to highly interdependent behavior, and allows the study of factors that control system-wide interdependence. Complexity theory and models developed in statistical mechanics (see e.g. Bak 1987 and Sachtjen et al 2000) can help explain how simple local rules create emergent system level behavior.

Third, banks react to changes in the environment - be these changes in policy or disruptions to the system's operation or changes in the behavior of other banks. Understanding how banks might react, and the impact of simultaneous reactions at the system level, greatly helps in evaluating risks and efficiencies of payment systems. While the incentives of banks may be analyzed individually in isolation, or when operating in a stipulated environment, their interaction in a system of banks with their own incentives necessitates a model. In modeling bank behavior, methodologies developed under reinforcement learning (Sutton and Barto 1998) and learning in games (Fudenberg and Levine 1998) may prove useful. As seen by the given example, already simple "intelligence" by agents can produce realistic behavior and add value to the analysis of payment systems. In the development of more realistic behavior for banks in settling payments, an important unanswered question is whether and what kind of bank behavior can be identified from empirical payment data.

References

Angelini, Paolo (1998). "An analysis of competitive externalities in gross settlement systems". J. Banking Finance 22, pp. 1–18.

Angelini, P., G. Maresca and D. Russo (1996). “Systemic risk in the netting system”, Journal of Banking and Finance, 20, pp. 853-68.

Albert, R.,H. Jeong, A.-L. Barabasi (1999). "Error and attack tolerance of complex networks". Nature 406, 378–382.

Bak, P., C. Tang, and K. Wiesenfeld (1987). "Self Organized Criticality: An explanation of 1/f noise", Physical Review Letters, 59:4:381-384.

Barabasi, A.-L. and R. Albert (1999), Emergence of scaling in random networks, Science 286, pp. 509–512.

Barto, A.G. and R.S. Sutton (1998). Reinforcement Learning: an introduction. MIT Press, Cambridge, Massachusetts.

Bech, M, Madsen, B and Natorp, L (2002). "Systemic Risk in the Danish Interbank Netting System". Danmarks Nationalbank working paper, No. 8.

Bech, M. and K. Soramäki (2002). Liquidity, gridlocks and bank failures in large value payment systems. in E-money and payment systems Review. Central Banking Publications.

Bech, M. and K. Soramäki (2005). "Systemic risk in a netting system revisited", in Leinonen, H. (ed.) "Liquidity, risks and speed in payment and settlement systems – a simulation approach". Bank of Finland Studies E:31.

Bech, M. and R. Garratt (2003). "The intraday liquidity management game", Journal of Economic Theory, vol. 109(2), pp. 198-219.

Bech, M. and R. Garratt (2006). “Illiquidity in the Interbank Payment System following Wide-Scale Disruptions”, Federal Reserve Bank of New York Staff Report 239.

Bedford, Paul, S. Millard, and J. Yang (2005). "Analysing the impact of operational incidents in large-value payment systems: a simulation approach", in Leinonen, H. (ed.) "Liquidity, risks and speed in payment and settlement systems – a simulation approach". Bank of Finland Studies E:31.

Bernanke, B. (1990). “Clearing and Settlement during the Crash”, Review of Financial Studies, 1990, v. 3, iss. 1, pp. 133-51.

Beyeler, Walter, M. Bech, R. Glass, and K. Soramäki (2006). "Congestion and Cascades in Payment Systems". Federal Reserve Bank of New York Staff Report No. 259.

Bollobas, B. (1985), Random Graphs, Academic Press, London, pp. 123-136.

Bonacich, P. (1972). "Factoring and weighting approaches to status scores and clique identification." Journal of Mathematical Sociology. 2:113-120.

Borgatti, S. P. (2005). "Centrality and network flow". Social Networks 27, pp. 55–71.

Brin, S. and L. Page (1998). "The anatomy of a large-scale hypertextual Web search engine". Proceedings of the seventh international conference on World Wide Web 7, 107-117.

Brown, G.W. (1951) "Iterative Solutions of Games by Fictitious Play" In Activity Analysis of Production and Allocation, T.C. Koopmans (Ed.), New York: Wiley.

Callaway, D., M. Newman, S. Strogatz, and D. Watts (2000). "Network robustness and fragility: Percolation on random graphs", Phys. Rev. Lett. 85, 5468-5471.

Crucitti, P., V. Latora, M. Marchiori, A. Rapisarda (2004a). "Error and attack tolerance of complex networks", Physica A 340, pp. 380–394.

Crucitti, P. V. Latora and M. Marchiori (2004b). "A model for cascading failures in complex networks" Phys. Rev. E 69, 045104.

Devriese, J. and J. Mitchell (2006). "Liquidity Risk in Securities Settlement", Journal of Banking and Finance, v. 30, iss. 6, pp. 1807-1834.

Dorogovtsev, S.N., J.F.F. Mendes, A.N. Samukhin (2001). "Giant strongly connected component of directed networks", Phys. Rev. E 64.

Dorogovtsev S.N. and J.F.F Mendes (2003). "Evolution of Networks". Oxford University Press.

Erdõs, P., and A. Rényi (1959), "On random graphs", Publicationes Mathematicae 6, pp. 290-297.

Freeman, L.C. (1978). "Centrality in social networks: I. Conceptual clarification." Social Networks, 1:215-39.

Fudenberg, D. and D.K. Levine (1998). The Theory of Learning in Games. MIT Press. Cambridge, Massachusetts.

Galbiati, M. and K. Soramäki (2007). " A competitive multi-agent model of interbank payment systems". forthcoming.

Hellqvist, M. and J. Koskinen (2005). "Stress testing securities clearing and settlement systems using simulations", in Leinonen, H. (ed.) "Liquidity, risks and speed in payment and settlement systems – a simulation approach". Bank of Finland Studies E:31.

Humphrey, D.B. (1986). "Payments Finality and Risk of Settlement Failure", in Technology and the Regulation of Financial Markets: Securities, Futures, and Banking. A. Saunders and L. J. White, eds. Heath: Lexington.

Inaoka, H., T. Ninomiya, K. Taniguchi, T. Shimizu, H. Takayasu (2004). "Fractal Network derived from banking transaction—An analysis of network structures formed by financial institutions". Bank of Japan Working papers No. 04-E-04.

Jensen, H. J. (1998). "Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems". Cambridge University Press, ISBN 0-521-48371-9

Kinney, R., P. Crucitti, R. Albert and V. Latora (2005). "Modelling cascading failures in the North American power grid". Eur. Phys. J. B 46, 101–107.

Kobayakawa, S. (1997). "The comparative analysis of settlement systems", Centre for Economic Policy Research Discussion Paper No. 1667.

Kuussaari, H (1996). "Systemic Risk in the Finnish Payment System: An Empirical Investigation". Discussion Paper 3, Bank of Finland.

Lacker, J.M. (2004). "Payment system disruptions and the federal reserve following September 11, 2001", Journal of Monetary Economics, Vol. 51, No. 5, pp. 935-65.

Leinonen, H. (2005). (ed.) "Liquidity, risks and speed in payment and settlement systems – a simulation approach". Bank of Finland Studies E:31.

Mazars, Emmanuel and Guy Woelfel (2005). "Analysis, by simulation, of the impact of a technical default of a payment system participant", in Leinonen, H. (ed.) "Liquidity, risks and speed in payment and settlement systems – a simulation approach". Bank of Finland Studies E:31.

McAndrews, J. and S.M. Potter (2002). "The Liquidity Effects of the Events of September 11, 2001", Federal Reserve Bank of New York Economic Policy Review, Vol. 8, No. 2, pp. 59-79.

McAndrews, J and S. Rajan (2000). "The timing and funding of Fedwire funds transfers", Federal Reserve Bank New York Econ. Pol. Rev. 6, pp. 17–32.

Moore, C., and M.E J. Newman (2000). "Exact solution of site and bond percolation on small-world networks". Physical Review E 62: 7059-7065.

Newman, M.E.J. (2003). "The Structure and Function of Complex Networks", SIAM Review, Vol. 45, No. 2, pp. 167—256.

Newman, M.E.J. (2005). "Power laws, Pareto distributions and Zipf's law". Contemporary Physics, Vol. 46, No. 5, 323 – 351.

Northcott, C.A. (2002) Estimating Settlement Risk and the Potential Contagion in Canada's Automated Clearing Settlement System. Working Paper 2002-41, Bank of Canada.

Sachtjen, M.L., B.A. Carrerras and V.E. Lynch (2000). "Disturbances in a power transmission system", Physical Review E, 61:5:4877-4882, 2000.

Soramäki K, M.L. Bech, J., R.J. Glass and W.E. Beyeler (2006). "The Topology of Interbank Payment Flows". Physica A. forthcoming.

Watkins, C. and Dayan, P. (1992). "Technical note: Q-learning". Machine Learning, 8(3/4), pp. 279-292.

Watts, D.J. (2002). "A simple model of global cascades on random networks". PNAS, April 30, 2002, vol. 99, no. 9, 5766-5771.