PRELIMINARY DRAFT

Systemically Important Accounts, Network Topology and Contagion in ARTIS

Michael Boss, Gerald Krenn, Valentina Metz, Claus Puhr, Stefan W. Schmitz*[1]

The purpose of this study is the investigation of the relevance of network topology for the stability of payment systems in the face of operational shocks. The analysis is based on a large number of simulations of the Austrian large value payment system ARTIS that quantify the contagion impact of operational shocks at participants' sites. The analysis uncovers that only few accounts are systemically important. We also find that network indicators at the node level can have some explanatory power. Their explanatory power is higher when the analysis focuses on the contagion measured by the number of banks with unsettled payments than in the case of the measure based on the value of unsettled payments. It is lower though than that of the more traditional measures of node activity (value and volume of payments). At this stage, network indicators at the network level seem to be of limited use for stability analysis.

JEL: E50, G10.

1 Introduction

Recent work on the stability of banking systems suggested a systematic relationship between network topology, system stability and contagion.[2] Similarly, a recent study conjectures that network topology might be relevant for the stability characteristics of payment systems.[3] In previous research we uncovered a large variation of the contagion impact across days, across banks, and across scenarios.[4] Here we investigate whether the position of the stricken account within the network explains its contagion impact and whether daily variations in network topology explain the variation of contagion across days.

In Section 2 we provide a brief motivation for studying network topology in network stability. In section 3 we present data on the network topology of ARTIS and compare it to the respective results for FedWire (the US large value payment system) and for the Austrian interbank market. Section 4 introduces the simulations. Based on the results we discuss the following questions: Which accounts cause contagion in the system and on what scale? How many are systemically important? Section 5 is devoted to the study of the following questions: First, do network indicators on the network level on the day of an operational failure relate to the contagion effects in the simulations? Second, do network indicators on the node level of the stricken participant on the day of the operational incident relate to the contagion effects in the simulations? Section 6 summarises the results.

[1] The authors thank DI Alfred Muigg, Mag. Wolfgang Draxler for providing data and valuable information. The views expressed in this paper are those of the authors and do not necessarily reflect those of OeNB and the Eurosystem. * Corresponding Author (stefan.schmitz@oenb.at).

[2] Boss et al. 2004.

[3] Soramäki et al. 2007.

[4] Schmitz, Puhr 2007.

2 Fundamentals of Network Topology and Network Stability

Many networks in the real world (e.g. Internet, WWW, large value payment systems such as FedWire and BOJ-NET, the Austrian interbank market) are scale-free networks. Their degree distribution follows a power law $P(k) \sim k^{-\gamma}$, i.e. the probability that a node has k degrees is $k^{-\gamma}$. A few nodes have a large number of links, while most nodes have only a few links. The network characteristics of scale-free networks are independent of the number of nodes and links. They are robust with respect to random node removal, but disintegrate quickly in cases of a targeted attack (the most highly connected nodes are removed at each step). Random networks constitute a different class of networks. They are characterised by a homogenous network structure, i.e. all nodes have a similar number of links. Random networks are less robust against random node removal, but are more stable with respect to targeted attacks than scale-free networks.

Albert, Jeong and Barabasi (1999, 2000) study the robustness of the World Wide Web (a subset of the WWW with 325 729 nodes and an average degree k 3.93) and the Internet (at the interdomain level with 6 209 nodes and k 4.59). In a stepwise procedure they remove a fraction of the nodes and links from the network. The node removals lead to the disappearance of all links to and from the removed nodes and to the decrease of the connectivity of the network. Some shortest paths between nodes become no longer available; some clusters of nodes that used to connect to the rest of the network get disconnected. In the case of random node removal a shock is simulated by removing a random sample of nodes. In the case of targeted attacks a shock is simulated by removing the most highly connected nodes in the network. They find that the size of the largest cluster of nodes in the WWW and the Internet decreases very slowly under random node removal, but rapidly under targeted attacks. Under the former the networks disintegrate when about 60 percent (WWW) and 80 percent (Internet) of all nodes are removed. Under the latter the networks break down already after the removal of as few as about 0.07 percent (WWW) and 0.03 percent (Internet) of all nodes, respectively. The authors explain the robustness results by the scale-free characteristics of the networks as most nodes have few links. As a consequence, random node removal is likely to hit lowly connected nodes with little implications for the connectivity of the entire network. The heterogeneity of the nodes and their distribution are also the reason for the low robustness with respect to targeted node removal. Already after a few rounds of removals most of the highly connected nodes that connect clusters of lowly connected nodes have disappeared and the network disintegrates.

How relevant are these results for the study of the stability of large value payment system with respect to operational problems at individual participants?

In Albert et al. the stability of the network is conceptualised as the connectivity of the remaining nodes and measured by the size of the largest cluster in the network and the average path length of the network. As the physical network structure of ARTIS is that of a complete network (Participants do not have to submit payments to each other via hubs; they can do so via direct links.), connectivity is not the relevant conceptualisation of stability. The stability problem is not that bank A cannot make a payment to bank C because the two are not linked anymore. The problem is that bank A might not have the liquidity to make the payment. As connectivity relates to the flow of liquidity in the system and the liquidity flows through hubs are higher than that through peripheral nodes, it plays an indirect role for the analysis of stability. Therefore, our measures of the contagion impact of shocks focus on the impact of the shock on the flow of liquidity (i.e. on the number of accounts with unsettled payment and on the value of unsettled payments) rather than on the disintegration of the network.

3 The Network Topology of ARTIS

The definition of the network under investigation is not trivial in empirical network analysis. In the analysis of network topology we focus on the Giant Strongly Connected Component (GSCC) of ARTIS.[5] The GSCC is the largest component of the network in which all nodes connect to each other via directed paths (i.e. without passing any node or link more than once). We have chosen this definition of the network for two reasons: first, ARTIS contains a comparatively large number of accounts which are not related to financial stability (i.e. small charities, offset accounts of OeNB's cash distribution subsidiary) and which are not active on most of the days in the sample. Second, we want to ensure the comparability of our data with that reported for FedWire in Soramäki et al. (2006) which refers to the GSCC.

Table 1: Network topology indicators (network level) in ARTIS (16 November 2005 to 16 November 2007) and in FedWire (Q1 2004) (averaged across days; network definition: GSCC)

	FedWire	ARTIS				
	Mean	**Mean**	**Median**	**Min**	**Max**	**Stdv**
Payments						
Volume	436 000	15 380	15 436	9 786	25 000	2 019
Value (EUR bn)	1 068	48.5	46.9	22.6	84.9	10.6
Average (EUR mn)	2.55	3.2	3.0	1.9	5.9	0.7
Connectivity						
Connectivity (%)	0.300	7.9	7.9	5.9	9.9	0.8
Distance Measures						
Avg. Path Length	2.6	2.4	2.4	2.2	2.6	0.080
Diameter	6.6	4.4	4	4	5	0.5
Others						
Clustering (%)	53.0	58.3	58.3	51.0	63.7	2.3
Average Degree	15.2	15.6	15.5	14.2	17.8	0.6
Betweenness Cent. (%)		0.8	0.8	0.6	0.9	0.1
Dissimilarity Index		0.47	0.47	0.39	0.60	0.03

Source: Own calculations (ARTIS) and Soramäki et al. 2006 (FedWire). Value and average value for FedWire are converted into Euro based on the USD/EUR exchange rate of 31 March 2004 of 1.21730.

In ARTIS the average volume of transactions per day is 15 380. The average value of transactions per day comes to EUR 48.5 billion. The average transaction size amounts to 3.2 million EUR. The size of the network is defined by the number of nodes n. On

[5] Mathematical definitions of the network indicators see the Appendix in Schmitz, Puhr (2007) and Zhou (2006). For comparable data on the network of all active accounts see Schmitz, Puhr (2007). For a description of the Austrian banking system see OeNB and FMA (2004) The Austrian Financial Markets, Vienna, pp. 50 55.

average there are 133.2 accounts in the GSCC during the sample period of which 63 are in the GSCC on all days. The active nodes are linked by an average of 1 376.1 directed links (m).[6] The connectivity p of the network is captured by the number of actual directed links relative to the number of possible directed links. Connectivity p averages 7.9 percent.

An indicator of the distance between nodes is the lowest possible number of links that connect each node with each other in the GSCC. It is referred to as shortest path length. We calculate the average shortest path length for each originating node by averaging across terminating nodes and then averaging across originating nodes to derive the average path length l of the entire network. Across days this value equals 2.4. This means that it takes only slightly more than two links to reach any terminating node in the network from any originating node in the network on average. Hence, the network is compact. This is mostly due to the fact that almost all active nodes are linked to the largest banks. This network structure is quite stable across days, as the standard deviation is low. The maximum path length across nodes is defined as diameter D. It is calculated by maximising across maximum path lengths which corresponds to picking an originating node at the very fringe of the network and counting the lowest possible number of links to the terminating node that is furthest away from it and leads to a value of 4.4 links.

How well are the nodes connected to the each other in the network? This is captured by the average degree k of the network which is calculated by summing across all (undirected) links originating from each node and than averaging across nodes.[7] Averaged also across days, it amounts to 15.6 in the ARTIS system. Pick a node in the GSCC on a random day in the sample period and it can be expected to have 15.6 links originating from (or terminating) at it. However, the most active nodes have a much larger number of links originating and terminating at them. The maximum out-degree averages 76 across days, so that the most active node on each day has about five times as many links originating from it than the average node. The maximum in-degree (90) is similarly much higher as the average degree.[8] The clustering coefficient provides a measure of the average connectivity of the neighbours of all nodes in the GSCC. On average about 58 percent of the neighbours of each node are also linked. Betweenness centrality measures how many shortest paths through the GSCC pass through the average node. The value of 8 percent is quite low and stems from the centrality of a few nodes with high betweenness centrality and a large number of nodes with low values. The dissimilarity index captures the relative viewpoints of the network from two neighbouring nodes. If the network looks very similar from both nodes, the dissimilarity index is small. In the GSCC it amounts to 0.47 which implies that on average the perspectives of the GSCC differ substantially from any two neighbouring nodes. A lot of nodes link that otherwise do not share many network characteristics. We interpret that as further evidence that many of the nodes connect to the largest nodes at the centre of the network.

---

[6] The average number of nodes in ARTIS active on every day was 209.8 and the number of directed links among them was 1 637.5.

[7] The out degree refers to the number of links originating at the node while the in degree is based on to the number of links terminating at the node. Across the network the average out and in degree are equal to m/n.

[8] For the network of all nodes active on every day the maximum out degree is 102 and the maximum in degree is 142.

How do these values compare to the results for FedWire? When comparing the two networks one must bear in mind that the FedWire data refers to Q1 2004. In particular value and volume in FedWire have certainly grown since then. The comparison between a small and a large network can yield interesting insights into the structure of payment systems. The average number of nodes in the GSCC of FedWire (n 5 086) is about 38 times that in ARTIS which implies that the number of possible directed links in FedWire is 1 469 times higher than in ARTIS. But the average number of directed links (m 76 614) is only about 55 times that in ARTIS so that connectivity should be lower in FedWire by a factor of about 26 (1 469 over 55). That is exactly the ratio between connectivity p in ARTIS (7.9 percent) and that in FedWire (0.3 percent). A conjecture based on this observation is that the number of possible directed links grows exponentially in payment systems, but the number of actual directed links only proportionally. The distance measures (average path length 2.6 vs. 2.4 and diameter 6.6 vs. 4.4), however, seem to be quite independent of size, like in other small-world networks.[9] The high clustering coefficients in both networks (on average 53 vs. 58 percent of the direct neighbours of each node are also linked) corroborate this finding. Also the average degrees of both networks are very similar (15.6 vs. 15.2).

Comparisons across networks are often based on the degree distribution. In scale-free networks it follows a Yule-Simon (or Power law) distribution $P(x) \sim k^{-\gamma}$ for degree values above a certain threshold. Many real world networks are said to follow a Power Law. The first indicator of the prevalence of the Power law is that the histogram of the degree distribution (on logarithmic scales) is a straight line with slope $-\gamma$, whereby in many real networks $-2 > -\gamma > -3$. The coefficient γ is estimated by a maximum likelihood estimator (e.g. Newman 2003). The respective value in Soramäki et al. (2005) is 2.11 for $k > 10$ for FedWire and that in Inaoka et al. (2002) is 2.3 for $k > 20$ for the BOJ-Net. Boss et al. report γ for the in-degree, the out-degree, and the degree distribution separately as 1.7, 3.1, and 2.0, respectively, for $k > 40$. For our monthly network[10] (degree range 1 to 1925 for the nodes in the GSCC over a period of 20 days) the histogram seems to indicate a Power law distribution with $\hat{\gamma}_{ML} = 1.4$ for $k > 10$ (see Graph 1a, left panel). However, Newman (2003) argues that the plot of the cumulative distribution function (cdf, on logarithmic scales) must also be a straight line with slope $-\gamma + 1$. Newman argues that the cdf plot is superior to the histogram, because it does preserve all the information in the data rather than throw out information by binning. In addition, it avoids the problem of noise in the tails that emerges from binning. We plot the cdf for the monthly network in the right panel of Graph 1a. The cdf obviously is not a straight line and we reject the Power Law hypothesis for the ARTIS network.

[9] In a small world network most nodes can be reached from every other by a small number of hops or steps, although connectivity is low and most nodes are not neighbors.

[10] We conducted the same exercise for the daily, the quarterly, and the semi annual networks with the same results.

Chart 1a: Histogram and (reverted) cumulated distribution function (on logarithmic scales) of the degree distribution in the monthly network in ARTIS (GSCC)

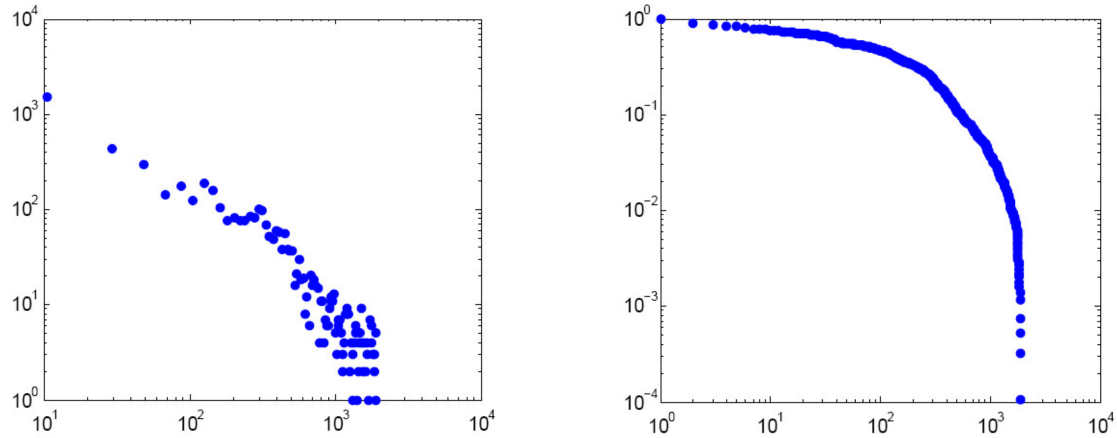

Left hand panel: y axis number of nodes, x axis *k* (degrees per node); Right hand panel: y axis $\hat{P}(x \geq k)$, x axis *k* (degrees per node); Source: Own calculations.

It is also interesting to compare the network indicators of the ARTIS system with the two network indicators of the Austrian interbank market presented in Boss et al. (2004) (data covers 2000 to 2003). Interbank market transactions can be interpreted as a subset of the transactions in ARTIS as they are settled through ARTIS. The authors find an average path length of 2.26±0.02 which is very close to the respective figure in table 1 of 2.3±0.05. That similarity arises from the fact that both, the interbank market and the payment system, are dominated by large banks. In both markets many banks cluster around their sectoral apex institutions.[11] However, the clustering coefficient is substantially higher in the ARTIS system than in the interbank network. Maintaining interbank relationships is costly so that banks have to balance the advantages of diversification with the costs of maintaining links. This is clearly not the case in the complete physical network of the large value payment system where the marginal costs of an additional link are zero. In addition, transactions in ARTIS are partly driven by customer payments (roughly 20 percent of total value). These reflect the network structure of real economic activity and that does not necessarily mirror that of the interbank market.

4 The Simulations: Methods, Data and Results

We conducted 31 311 simulations based on 63 different scenarios for 497 transaction days from 16 November 2005 to 16 November 2007 (excluding Austrian holidays) which

[11] Of the seven sectors the Raiffeisen credit cooperative, the Volksbanken credit cooperatives and the Savings Banks have a tiering structure. They account for about 80 percent of Austrian banks in terms of the number of credit institutions and for about 50 percent in terms of total assets (unconsolidated). In addition, there is no national automated clearing house in Austria and the Austrian banking system relies on correspondent banking relationships to settle a range of customer payments (e.g. credit transfers). The banks that operate in ARTIS do have direct access to the system based on their own in house systems. Although within sectors the IT solutions are often similar, there is no evidence that operational risk is correlated across individual banks within a sector.

yielded some 620 million simulated transactions.[12] These simulations were calculated with a self-implemented Matlab based software tool (inspired by Bank of Finland Payment System Simulator), which was tailored to ARTIS particularities. The tool recalculates the transactions of each day by adding incoming payments to and subtracting outgoing payments from the respective accounts of the participants. As transactions in the input data set provide time stamps, the simulator recalculates the balances of all participants of the system throughout the day depending on the institutional features of the system (e.g. settlement algorithm, queue release mechanism). The institutional features of the system that could not be accounted for in the simulator had to be mapped into the input data set. Since the tool cannot take into account behavioural reactions of system participants, they must be determined exogenously. First of all, other participants might want to stop submitting payments to the participant experiencing operational problems. In TARGET a stop sending rule applies, if a transfer account of a central bank in the system experiences an operational problem. In this case, no further payments are transferred to the stricken transfer account.[13] Payments to other participants are not affected. In cases of operational problems at another bank, ARTIS operators provided evidence that in all other cases participants continue to submit payments to participants that experience operational problems, even if the latter cannot submit payments themselves for many hours. Second, participants could react to operational incidents by increasing available collateral. Anecdotal evidence suggests that participants already hold large shares of their eligible assets at OeNB. Consequently, we assume that system participants are not increasing collateral for durations of operational incidents of up to one day. The simulations are based on actual liquidity data for the sample period. We interpret the sum of beginning of day balances on ARTIS accounts plus unencumbered eligible collateral held at OeNB as the binding liquidity constraint for banks. Third, the simulation algorithm takes into account debit authorisation by the bank for a number of other participants in ARTIS.[14]

The scenarios in Schmitz, Puhr (2007) were designed on the basis of the analysis of actual payment flows in ARTIS focusing on the most active accounts which also featured the highest risk concentration measures during the sample period.[15] This resulted in three scenarios: in the first, the most active transfer account[16] was shocked; in the second, the

[12] For more details on simulations, their motivation, and their design see Schmitz, Puhr (2007). The operation of ARTIS was discontinued after 16 November 2007, due to the introduction of TARGET2.

[13] Due to the operating procedures it actually takes about 40 minutes after the detection of the operational problem at the transfer account until a stop sending is imposed. The implementation of the rule in the simulation algorithm takes that small delay into account.

[14] Participant A can grant participant B a debit authorisation according to the Terms and Conditions Governing the OeNB's ARTIS system (§ 9). Debit authorisation is defined as the right of participant B to initiate (certain pre agreed) payments from the account of participant A. Debit authorisations are granted to a small number of participants for prearranged purposes (very frequent recurring standard operations) and cannot be interpreted as crises mitigation instrument available on short notice in the case of an operational incident.

[15] The measures employed were (1.) the value of liquidity concentrated at the nodes, (2.) the number and value of payments submitted and received (payment concentration channel), (3.) the Herfindahl index of concentration of payment flows (both based on the number of payments and the value of payments received and submitted) as well as (4.) the monthly network topology.

[16] Transfer accounts are ARTIS accounts held by other ESCB central banks at OeNB. All national TARGET components are directly linked by transfer accounts. All transactions to and from the respective country and Austria are routed via these transfer accounts.

most active bank account was assumed to experience operational problems; and in the third one, the three most active bank accounts were stressed simultaneously.

In this paper we run simulations for all 50 banks that are in the GSCC on all Austrian working days throughout the sample period and all 13 transfer accounts that form part of the system on all days in the sample period. We assume an operational incident that hits one account in each simulation. The operational incident is mapped into the simulation as the incapacitation of the participant to process outgoing payments, i.e. the inability to submit transactions for the whole day.[17] This assumption is extreme but plausible. Shorter outages of participants may lead to payment delays but not to unsettled payments, as shown in Schmitz, Puhr (2007).

The results are graphically depicted in the four panels of chart 1. In the upper left hand panel the number of contagious defaults per simulation (in terms of the number of banks with unsettled payments) is depicted on the x-axis; the numbers of simulations that yield x contagious defaults on the y-axis. It shows that about 27 percent of all simulations (8 604) do not lead to contagion at all. A further 26 percent (8 230) yield one contagious default and 16 percent (4 919) two. About 29 percent (5 456) lead to three to five contagious defaults and 17 percent (4 102) to more than five. Maximum contagious defaults across the 31 311 simulations amount to 33.

The time series of average contagious defaults (in terms of the number of banks with unsettled payments) per day is featured in the upper right hand panel. It is quite volatile with a standard deviation of about 25 percent of the mean. This motivates the investigation in section 5.1 where we study the question whether the variation of network topology across days can contribute to the explanation of the fluctuations of average contagious defaults per day.

[17] It is assumed that the resulting illiquidity of the participant is not interpreted as potential insolvency by other participants of the payment system and the financial system at large. In addition, ARTIS provides business continuity arrangements for participants. We tested their impact in Schmitz and Puhr (2007), but disregard them in this paper, as they are of little relevance for the interaction between network topology and contagion.

Chart 1

Simulation results

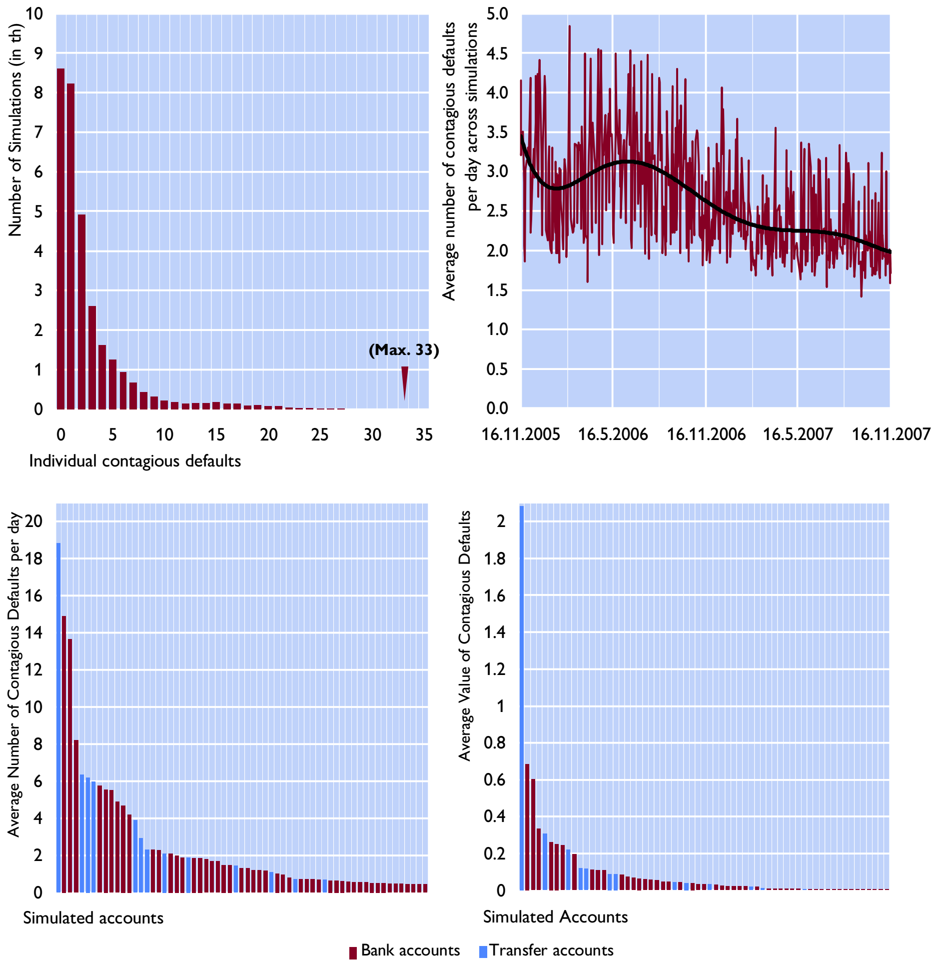

Source: OeNB.

The lower panels in chart 1 show the average contagious defaults per simulation (in terms of number of banks with unsettled payments, lower left hand panel) and the average value of unsettled payments due to contagious defaults (lower right hand panel) per simulation. We use this information to derive the set of systemically relevant accounts in the following way: As argued above, connectivity is not an adequate criterion to capture the systemic impact of an operational problem at one of the nodes in a large value payment system. Alternatively, we suggest defining a threshold based on the average contagion effect of an individual account. The latter can be measured in terms of the number of contagious defaults or in terms of the value of unsettled payments in the system. To some extent that threshold is arbitrary and depends on the risk aversion of the supervisory authority. If we set the threshold in terms of the number of contagious defaults at 1, i.e. only accounts that yield at least an average of one bank with unsettled payments due to contagious default across the sample period, we find that only 39 accounts in the GSCC are systemically relevant which includes eleven transfer accounts operated by central banks (lower left hand panel in chart 1). These 28 bank accounts amount to twelve percent of the average of 230 bank accounts in ARTIS (during the sample period) and to about 3 percent of the average of 850 banks in Austria. If we define the threshold in terms of the value of contagious defaults, i.e. only accounts that cause at least an average value of EUR 48.5 million of unsettled payments (or 0.1 percent of average value of transactions settled across days), we find that 24 accounts are systemically relevant (lower right hand panel in chart 1). Seven of these are transfer accounts and 17 bank accounts which account for about seven percent of the average of 230 bank accounts in ARTIS (during the sample period) and to about two percent of the average of 850 banks in Austria.

Given that transfer accounts do not hold any liquidity (i.e. the liquidity drain caused by their incapacitation is nil) and that the stop sending rule strongly reduces the liquidity sink effect, the strong contagion impact of transfer accounts is interesting. It indicates that payment concentration risk is more important for the contagion impact than liquidity concentration risk. TARGET2 operates on a Single Shared Platform without the highly contagious transfer accounts. This could increase the resilience of this critical infrastructure with respect to operational problems (though not necessarily at the platform level).

The results suggest that the supervision of operational risk in banks' payment processing/submission capacity could focus on a relatively small set of systemically relevant banks in Austria and on their business continuity arrangements.

Approximating a probability distribution across contagious defaults per simulation

In section 3 we showed that large value payment systems can have common network characteristics despite large differences in size. In order to provide an opportunity to compare the simulation results across large value payment systems, we estimate the relation between the number of simulations and the number of contagious defaults they cause (in terms of the number of banks with unsettled payments). Chart 1 (upper left panel) reveals that the number of simulations y that involve a certain number of contagion events x is a rather regularly declining function in x. In this context it seems natural to look for a simple parametric probability distribution describing the number of occurrences of contagion events in a simulation, given that contagion actually did occur. As such a distribution would attach positive probabilities to low probability high impact events, it could be applied in future simulation studies for the analysis of extreme events.

As candidate distributions, we considered discretised versions of the following continuous distributions: Exponential, Weibull and Gamma. These three distributions are defined on the set of non-negative numbers and have one (Exponential) or two (Weibull and Gamma) parameters. Discretising these distributions was accomplished in the following way: The probability of observing just one contagion event was set to the probability of observing the continuous distribution in the interval from zero to one; observing two contagion events was related to the interval from one to two; and so on. The maximum likelihood method was used for estimating the unknown parameters.

A graphical assessment of the adequacy of the estimated distributions shows that exponential distributions are not flexible enough in order to describe the observed numbers of contagion events. This is due to the fact that this distributional family only has a scale but no form parameter. A much better fit is achieved by the Weibull and Gamma distributions. When applying chi square tests for goodness of fit, however, it comes as no surprise that these distributions are rejected at any commonly used confidence level as we are dealing with a very large number of observations (22 707).[18] Nevertheless, it can be observed that the Weibull distribution delivers a smaller value of the chi square statistic than the Gamma, thus indicating a better fit of the former. We conclude that a reasonable choice for describing the probability that C, the number of contagion events in a simulation that actually show at least one contagion event, is equal to a positive integer n given by:

$$\mathrm{P}\{C = n\} = \mathrm{Wei}(n \mid a,b) - \mathrm{Wei}(n-1 \mid a,b) \quad \text{for all } \; n \geq 1,$$

where $\mathrm{Wei}(. \mid a,b)$ denotes the cumulative distribution function of a Weibull distribution with parameters a and b, defined by

$$\mathrm{Wei}(x \mid a,b) = 1 - \exp(-(x/a)^b) \quad \text{for all } \; \mathrm{x} \geq 0.$$

We thus approximate the distribution of the number of contagious defaults given that contagion actually occured by means of a discretised Weibull distribution with $\hat{a}_{ML} = 2.61$ and $\hat{b}_{ML} = 0.77$.

5 The Interaction between Network Topology and Stability in ARTIS

In this section we investigate whether the variation of network indicators at the network level across days (5.1) and at the node level across stricken accounts (5.2) explain the variation of contagion across days and across stricken accounts.

The selection of the appropriate measure of network topology is not trivial as the number of available indicators is large. At the network level we calculate 44 network indicators taking into account not only those in table 1 but also the directed and/or value/volume weighted and/or average/maximum values for selected indicators. Similarly, the number of available indicators at the node level comes to 71.

Boss et al. (2004) relate contagion in the interbank market to betweenness centrality at the node level, because this measure has a higher explanatory value than the alternative network indicators in their data set. They uncover a dented linear relationship. Banks with betweenness centrality measures $0 \leq C_B(h) \leq 2$ do not cause any contagious defaults. For $C_B(h) > 2$ they find a linear relationship with a slope of about 0.8.

Borgatti (2005) studies the selection of the appropriate centrality measure for various typologies of flow processes. He classifies flows along two dimensions: the characteristics of the route through the network and the characteristics of the transfer mode. The first dimension encompasses paths, trails, and walks. Paths are sequences of links and nodes in which neither links nor nodes are repeated (Shortest paths are a special case of paths.) Trails refer to sequences in which nodes but not links may be repeated. Walks are unconstrained sequences. The second dimension refers to the way in which the flowing good is passed on along the route from one node to another. While a disease can be

[18] Due to the large sample size even small deviations of the fitted values from the observed values lead to a formal rejection of the null hypothesis which reflects a common criticism to statistical tests (De Groot 1985).

passed on without implying the immediate cure of the carrier (Borgatti refers to this as parallel duplication), liquidity is transferred so that the initial holder has to part with it (referred to as transfer). What does that imply for the flow of liquidity in ARTIS? In a physically complete network banks do not have to make payments to other banks via third parties. They transfer directly to the ultimate receiver. However, the flow of liquidity does not stop there. It can be transferred to any other node in the network (including the submitter of the first payment). Where liquidity ultimately ends up, is beyond the control (and interest) of the initial submitter of a payment. That implies that liquidity flow follows a walk rather than a path or a trail. Given that betweenness centrality is based on the share of all shortest paths through a node, it is not a good measure of centrality in the study of liquidity flows. Degree centrality is more suitable for this purpose.

We present our results in terms of four network indicators for three reasons: First, we believe that given the nature of liquidity flows degree centrality is the appropriate measure; second, we want to ensure a high degree of comparability of our results with other papers that use different network indicators (i.e. betweenness centrality), third, we want to investigate whether network indicators in general add value to the more traditional measure used in comparable simulation studies (i.e. the size of the individual node in terms of value and volume of transactions). Therefore we focus on the measures value and volume as well as on the network indicators degree, average path length, betweenness centrality and dissimilarity index in each of the following two subsections.

5.1 Network Level

In chart 2 we depict the daily values for the value (left hand panel) and the volume of all payments (right hand panel) submitted to ARTIS on the y-axis and the number of contagious defaults (in terms of the number of banks with unsettled payments – daily averages across scenarios) per day on the x-axis. The variation of value explains 2 percent and the variation of volume accounts for 8 percent of the variation of the contagion impact per day.

Chart 2

Value and volume (network level) per day versus average number of contagious defaults per day

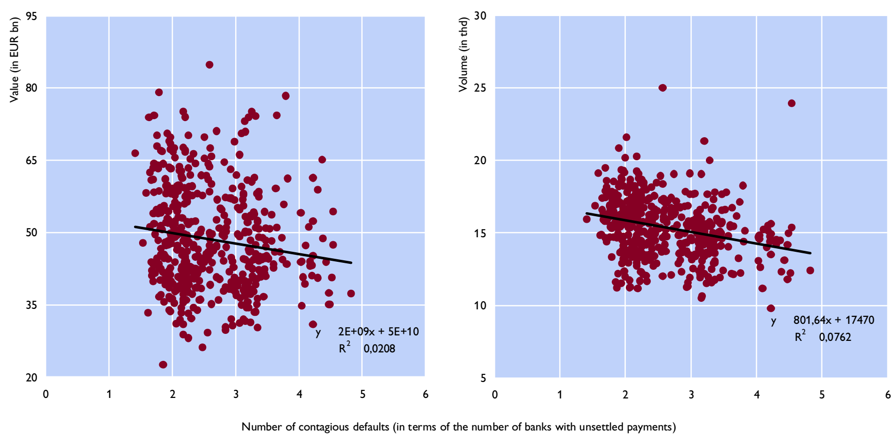

Source: OeNB.

The explanatory value of the variables value and volume is low. Do network indicators perform any better? In chart 3 we look at the following indicators (in and out, unweighted, undirected): degree, average path length, betweenness centrality and dissimilarity index. Similarly to chart 2 the daily number of contagious defaults (in terms of banks with unsettled payments – daily averages across scenarios) is depicted on the x-axis and the daily values of the respective network indicator are shown on the y-axis in each panel.

Chart 3

Selected network indicators (network level) per day versus average number of contagious defaults per day

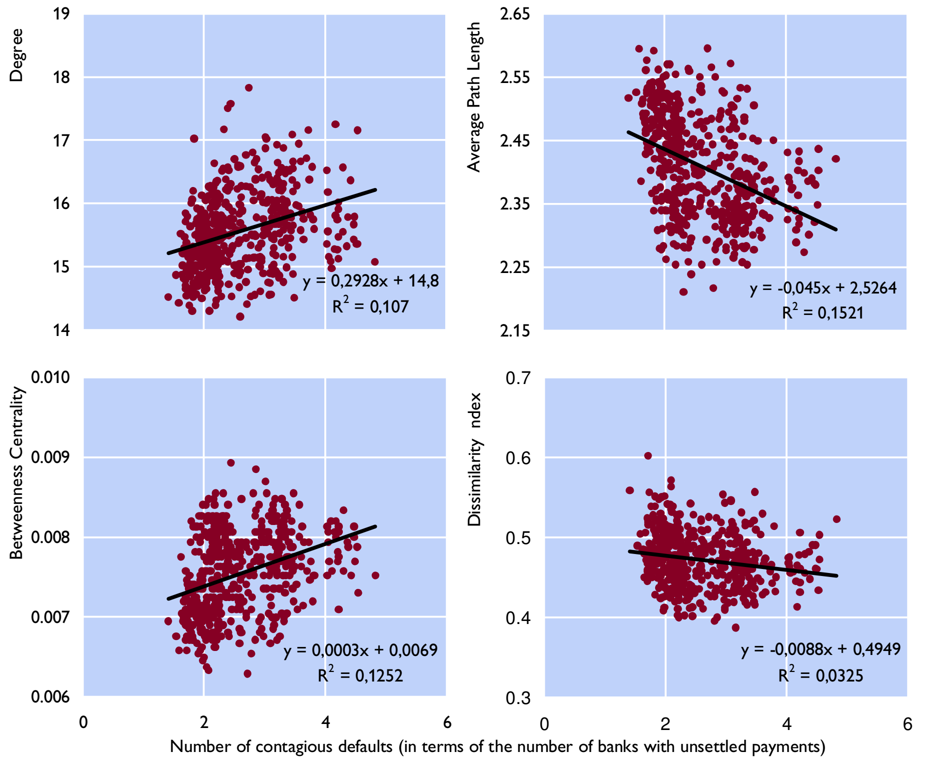

Source: OeNB.

The average path length (15 percent) and betweenness centrality (13 percent) have the highest explanatory values. The daily variation in degree accounts for 10 percent of the variation in contagion and that of the dissimilarity index for only 3 percent. Although the explanatory power of three of the network indicators is higher than that of value and volume, the levels are still low. The highest explanatory power of any of the remaining 39 indicators is 15.4 percent (average number-weighted clustering coefficient), while a number of indicators have no explanatory power at all. We conclude that daily variations in network structure are of limited use in the stability analysis of ARTIS. However, that does not preclude that structural differences across networks might influence their relative resilience. But as shown above, even large value payment systems which display considerable differences in size share notable structural commonalities.

5.2 Node Level

In this subsection we study the large dispersion of contagion effects caused by different nodes (see lower panels in chart 1). Do the different positions of the nodes (that experience the operational shock) in the network account for this variation? In chart 4 we plot the value and volume of payments of the stricken node in each simulation against its contagion effect in terms of number of contagious defaults (in terms of the number of banks with unsettled payments) i.e. each sub-graph contains 31 311 data point. In addition, the data points of the three most active banks (Bank A, B, and C) and of the most active transfer account (Transfer Account 1) are coloured (see chart 4 legend) while those of all other Bank Accounts and of all other Transfer Accounts are black and grey, respectively. The variations of value and volume across simulations explain 73 percent and 68 percent of the variation of the contagion impact across simulations. The slopes have the expected signs: more active nodes cause more contagion. The differentiation among simulations according to the shocked account reveals a pronounced grouping in both panels. In the right hand panel it also points to structural differences in contagion impact not accounted for by variations in volume. Transfer Account 1 and Bank B tend to group below the regression line (i.e. they causes more contagion than estimated by their volumes of transactions) and Banks A and C to group above the regression line (i.e. they cause less contagion than estimated by their volumes of transactions).

Chart 4

Value and volume (node level) per stricken account versus number of contagious defaults per simulation

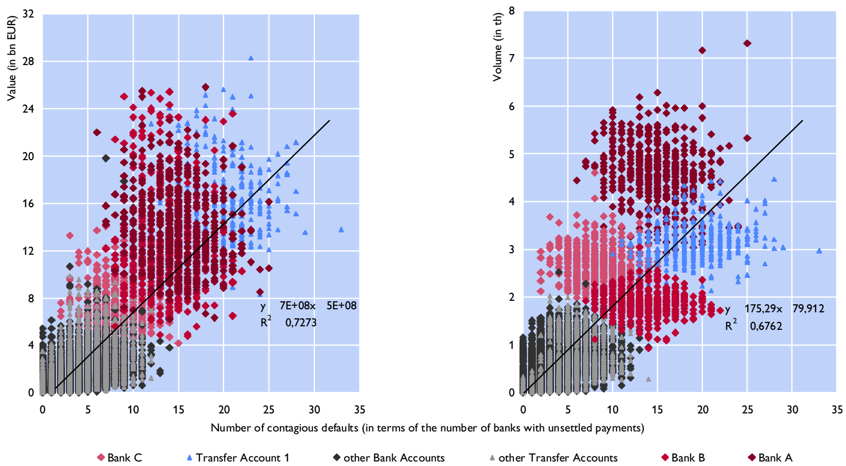

Source: OeNB.

In chart 5 we plot four network indicators (degree, average path length, betweenness centrality and dissimilarity index) of each stricken node against its contagion effect in terms of number of contagious defaults (i.e. each sub-graph contains 31 311 data points). In addition, the data points of Banks A, B, and C and Transfer Account 1 are differentiated in the same way as in chart 4. The explanatory values of all four network indicators are quite high; the simplest measure degree yields an R^2 of 64 percent, variations in average path length across simulations account for 59 percent of the variation of the number of contagious defaults across simulations. The more complex measures betweenness centrality and dissimilarity index yield R^2s of 52 and 62 percent, respectively. These values are in the order of magnitude of the reported interaction between betweenness centrality and contagious defaults for the Austrian interbank market (Boss et al. 2004). The signs of the slopes are in line with expectations: simulations in which more active and more central nodes are shocked feature a higher contagion impact. The remaining 65 network indicators yield explanatory values between nil (number-weighted average path length based on payments received) and 77 percent (relative volume of payments received). The results demonstrate that network indicators at the node level can indeed explain large parts of the variation in contagion across stricken accounts. However, they seem to add little to the high explanatory values of the traditional measures of activity (value and volume). Furthermore, the large set of available indicators and the huge differences in their explanatory values pose the problem of data mining. The differentiation according to the stricken account confirms the pronounced grouping evident also in chart 4. In all four panels simulations based on Transfer Account 1 cluster at the right hand side of the regression line, while those based on Bank C and to a lesser extent those of Bank A lie to its left. This finding points at structural differences in contagion impact which are not accounted for by measures of activity or network indicators and which warrant further research.

Chart 5

Network indicators (node level) per stricken account versus number of contagious defaults per simulation

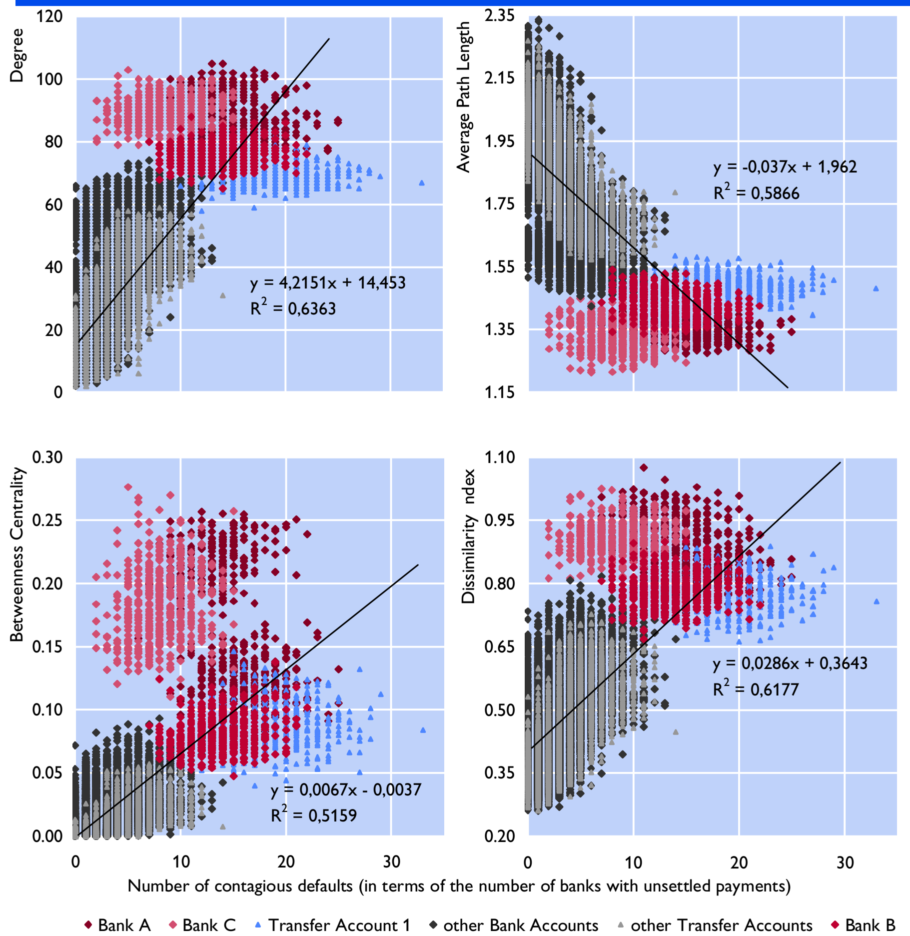

Source: OeNB.

We also investigate the interaction between network topology and network stability for another measure of contagion, i.e. the value of unsettled payments. Again we start with the analysis of the explanatory of node size (value of value and volume of payments originating at the node) (chart 6). Variations in value explain 54 percent and volume 39 percent of the variation in contagion. Both values are lower than the respective results in chart 4.

Chart 6

Value and volume (node level) per stricken account versus value of contagious defaults per simulation

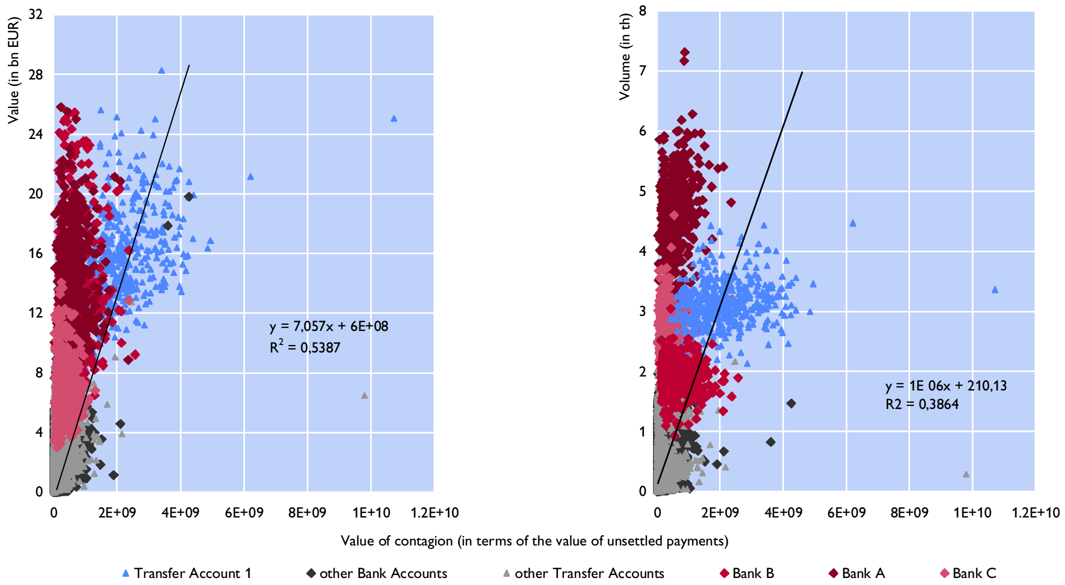

Source: OeNB.

How well do the network indicators at the node level fare in comparison? The explanatory values are similar for the four network indicators (degree 28 percent, average path length 25 percent, betweenness centrality 24 percent and dissimilarity index 29 percent, chart 7) and they are considerably lower than the respective values for the measures of size in chart 6. If contagion is measured by the value of unsettled payments, network indicators are clearly dominated by the traditional measures of size. However, the grouping of contagious defaults according to the three most active bank accounts and the most active transfer account are also apparent in charts 6 and 7. Comparing the results for the two measures of contagion, number of banks with unsettled payments (charts 4 and 5) versus value of unsettled payments (chart 6 and 7), reveals that contagion under the latter measure is much harder to explain by the more traditional variables (value and volume of payments) and by network indicators. But relatively speaking network indicators do even worse. In future work, we will focus on the investigation of the variations in the value of contagion in a multivariate setting in which we combine control variables (e.g. beginning of day liquidity at individual nodes) with network topology indicators at the network and at the node level.

Chart 7

Network indicators (node level) per stricken account versus value of contagion per simulation

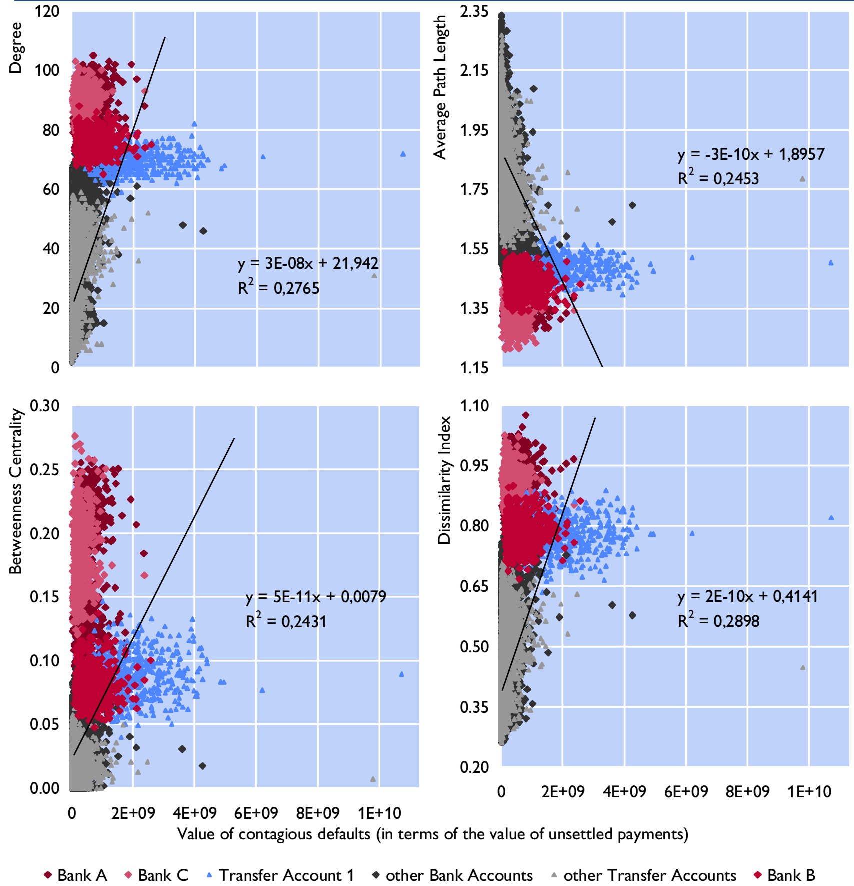

Source: OeNB.

In order to corroborate our finding that network indicators at the node level do not add much value to stability analysis, we present the correlations between the traditional measures of activity (value and volume) and selected network indicators in table 2. The data reveals that various indicators of centrality (average path length, degree, connectivity, betweenness centrality and dissimilarity index) are highly correlated with value and volume.

Table 2: Correlations between network indicators (node level)

	Volume	Value	Avg. PL	Degree	Conn.	Clust.	Btw. C.	Dissim.
Volume	100%	89%	77%	84%	83%	57%	89%	85%
Value		100%	70%	76%	75%	52%	77%	78%
Avg. PL			100%	96%	97%	62%	79%	85%
Degree				100%	99%	72%	85%	95%
Conn.					100%	72%	85%	93%
Clust.						100%	56%	78%
Btw. C.							100%	87%
Dissim.								100%

Source: OeNB. Average Path Length (Avg. PL), Connectivity (Conn.), Clustering Index (Clust.), Betweenness Centrality (Btw. C.), Dissimilarity Index (Dissim.).

The analysis suggests that network indicators provide little value added in the stability analysis of large value payment systems with respect to operational shocks at a participant. In future research we will extend the analysis from a univariate to a multivariate framework.

6 Summary

The analysis of the network indicators of ARTIS shows that the network is compact. This is mostly due to the fact that almost all active nodes are linked to a small number of accounts at the centre of the network (the largest banks and the most active transfer accounts). This network structure is quite stable across days. The comparison between the ARTIS system and the much larger FedWire network yields interesting insights into the relationship between size and structure of payment systems. The distance measures, the average degree, and the clustering coefficient seem to be independent of size, like in other small-world networks. A comparison of the network indicators of the ARTIS system with those of the Austrian interbank market reveals that the distance measures are very similar but the clustering coefficients differ substantially. That similarity arises from the fact that the interbank market is dominated by a few large nodes in the centre of the network, too.

We conducted 31 311 simulations based on 63 different scenarios for 497 transaction days from 16 November 2005 to 16 November 2007 (excluding Austrian holidays). Although the scenarios focus only on the banks and on the transfer accounts in the GSCC on all days, more than a quarter of all simulations do not lead to contagion (in terms of the number of banks with unsettled payments) at all, and two fifth yield one or two contagious defaults. Based on two conservative thresholds of contagion impact we find that only a very small number of accounts are systemically important. If we regard only accounts that yield at least an average of one contagious default across the sample period as systemically important, we find that only 28 bank accounts are systemically relevant, but almost all transfer accounts operated by central banks. If we define systemic relevance as contagion impact of at least 0.1 percent of the average value of transactions settled across days, we find that 17 bank accounts and seven transfer accounts are systemically relevant. In both cases only seven to twelve percent of all bank accounts in ARTIS and two to three percent of all Austrian banks are systemically relevant. The simulation results suggest that the ARTIS system is remarkable stable with respect to operational incidents at one of its participants. The strong contagion impact of the transfer accounts is an interesting feature revealed by the simulations and suggests that the removal of transfer accounts by the single shared platform in TARGET 2 can improve resilience relative to the old TARGET system.

The time series of average contagious defaults per day is quite volatile. We find that the variation of network structure across days does not contribute much to the explanation of the variation of contagion across days. At this stage, network indicators at the network level seem to be of limited use for stability analysis.

Network indicators at the node level can have explanatory power. In the simulations some of them are correlated with the contagion impact of an operational shock to a node. Their explanatory power is higher when the analysis focuses on the contagion measured by the number of banks with unsettled payments than in the case of the measure based on the value of unsettled payments. It is questionable at this stage that they contain much additional information compared to value and volume which traditionally were the focus of stability analysis in simulation studies of operational risk in large value payment systems. Furthermore, the large number of available network indicators at the node level and the huge differences in their explanatory power pose the problem of data mining. In future research we plan to explore the large data set compiled in the simulations to investigate the explanatory power of network indicators at the network and at the node level in a multivariate framework which allows controlling for other explanatory variables such as beginning of day liquidity at the network and at the node level.

6 Bibliography

Albert, R., H. Jeong, A.-L. Barabasi (1999) Diameter of the World Wide Web, Nature 401, 130-131.

Albert, R., H. Jeong, A.-L. Barabasi (2000) Error and attack tolerance of complex networks, Nature 406, 378-381.

Albert, R., A.-L. Barabasi (2002) Statistical mechanics of complex networks, Reviews of Modern Physics 74, 47-97.

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

Boss, M., H. Elsinger, M. Summer, S. Thurner (2004) An empirical analysis of the network structure of the Austrian interbank market, OeNB Financial Stability Review 7, 77-87.

DeGroot, M.H. (1985) *Probability and Statistics*, Second Edition, Addison-Wesley. Reading, Massachusetts.

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.

Latzer, M., Schmitz, S. W. (eds.) (2002) *Carl Menger and the Evolution of Payment Systems. From Barter to Electronic Money*, Edward Elgar, Cheltenham.

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

Newman, M. E. J. (2003) The structure and function of complex networks, (available at http://arxiv.org/abs/cond-mat/0303516).

Newman M. E. J. (2005) Power Laws, Pareto Distributions, and Zipf's Law, Contemporary Physics 46, 323-351.

Oesterreichische Nationalbank and Finanzmarktaufsichtsbehörde (2004) The Austrian Financial Markets, Vienna.

Schmitz, S. W., C. Puhr (2006) Liquidity, Risk Concentration and Network Structure in the Austrian Large Value Payment System, Available at SSRN: http://ssrn.com/abstract=954117.

Schmitz, S. W., Puhr, C. (2007) Risk concentration, network structure and the assessment of contagion in the Austrian large value system ARTIS, in: H. Leinonen (ed.), *Simulation studies of the liquidity needs, risks and efficiency in payment network*, Bank of Finland Scientific Monograph E:39, Helsinki,183-226.

Schmitz, S. W., G. E. Wood (eds.) (2006), *Institutional Change in the Payments System and Monetary Policy*, Routledge, London.

Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass, W. E. Beyeler (2006) The Topology of Interbank Payment Flows, Federal Reserve Bank of New York Staff Report, New York No. 243.

Soramäki, K., W. E. Beyeler, M. L. Bech, R. J. Glass (2007) New approaches for payment system simulation research, in: H. Leinonen (ed.) *Simulation studies of the liquidity needs, risks and efficiency in payment network*, Bank of Finland Scientific Monograph E:39, Helsinki, 15-40.

Zhou, H. (2003) Distance, dissimilarity index, and network community structure, *Physical Review* E 67, 061901, 1-8.