

Systemically Important Accounts, Network Topology and Contagion in ARTIS

This study investigates 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. It uncovers that only few accounts are systemically important. We also find that network indicators at the node level can have some explanatory power, which is higher when the analysis focuses on contagion measured by the number of banks with unsettled payments than on that measured by the value of unsettled payments. The explanatory power is, however, lower than that of the more traditional measures of node activity (value and volume) of payments. At this stage of our research, network indicators at the network level seem to be of limited use for stability analysis.

Michael Boss,
Gerald Krenn,
Valentina Metz,
Claus Pühr,
Stefan W. Schmitz¹

JEL classification: E50, G10

Keywords: Operational risk, payment systems, network topology, stress tests

1 Introduction

Recent work on the stability of banking systems suggested a systematic relationship between network topology, system stability and contagion (Boss et al., 2004). Similarly, Soramäki et al. (2007) conjectured that network topology might be relevant for the stability characteristics of payment systems. In previous research (Schmitz and Pühr, 2007), we uncovered a large variation of the contagion impact across days, banks as well as scenarios. Here, we investigate whether the position of the stricken account within the network explains its contagion impact and whether daily variation in network topology explains 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 the Austrian large-value payment system ARTIS and compare them with the respective results for the U.S. large-value payment system FedWire and for the Austrian interbank market. Section 4

introduces the simulations, and based on the results, we discuss the following questions: Which accounts cause contagion in the system and on what scale? How many accounts are systemically important? In section 5 we address the questions: Do network indicators at the network level on the day of an operational failure relate to the contagion effects in the simulations? And second, do network indicators at the node level of the stricken participant on the day of the operational incident relate to the contagion effects in the simulations? Section 6 summarizes the results.

2 Fundamentals of Network Topology and Network Stability

Many networks in the real world (e.g. the Internet, the World Wide Web, large-value payment systems, such as FedWire in the U.S.A. and BOJ-NET in Japan, 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

¹ Oesterreichische Nationalbank, Financial Markets Analysis and Surveillance Division; michael.boss@oebn.at, gerald.krenn@oebn.at, valentina.metz@oebn.at, claus.puhr@oebn.at, stefan.schmitz@oebn.at. The views expressed in this paper are those of the authors and do not necessarily represent the views of the OeNB and the Eurosystem. The authors thank Alfred Muigg and Wolfgang Draxler (both OeNB) for providing data and valuable information.

Refereed by:
Harry Leinonen,
Bank of Finland

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 case of a targeted attack, when the most highly connected nodes are removed step by step. Random networks constitute a different class of networks. They are characterized 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 et al. (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 inter-domain level with 6,209 nodes and $k=4.59$). They remove a fraction of the nodes and links from the network in a step-wise procedure. The node removals cause the disappearance of all links to and from the removed nodes and reduce the network's connectivity. 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 and, in the case of targeted attacks, by removing the most highly connected nodes in the network. Albert et al. (1999, 2000) 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% (WWW) and 80% (Internet) of all nodes are removed. Under the latter, the networks break down after the removal of as few

as about 0.07% (WWW) and 0.03% (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 networks' low robustness against targeted node removal. Even after just a few rounds of removals, most of the highly connected nodes that link 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 systems with respect to operational problems at individual participants?

In Albert et al. (1999, 2000) the stability of the network is conceptualized 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 may submit payments to each other via direct links rather than via hubs), connectivity is not a useful conceptualization of stability. The stability problem is not that Bank A cannot make a payment to Bank C because of a broken link, but that Bank A might not have adequate liquidity. As connectivity relates to the flow of liquidity in the system and the liquidity flows through hubs are higher than those 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 effects shocks have on the flow of liquidity (i.e. number of accounts with unset-

Table 1

Network Topology Indicators (Network Level) in ARTIS (November 16, 2005 to November 16, 2007) and in FedWire (2004/Q1) (Averaged across Days; Network Definition: GSCC)

	FedWire		ARTIS			
	Mean	Mean	Median	Min	Max	Stdv
Payments						
Daily volume (number of transactions)	436,000	15,380	15,436	9,786	25,000	2,019
Daily value (EUR billion)	1,068	48.5	46.9	22.6	84.9	10.6
Average value per transaction (EUR million)	2.55	3.2	3	1.9	5.9	0.7
Connectivity measures						
Connectivity (%)	0.3	7.9	7.9	5.9	9.9	0.8
Distance measures						
Average path length	2.6	2.4	2.4	2.2	2.6	0.08
Diameter	6.6	4.4	4	4	5	0.5
Other measures						
Clustering (%)	53	58.3	58.3	51	63.7	2.3
Average degree	15.2	15.6	15.5	14.2	17.8	0.6
Betweenness centrality (%)	–	0.8	0.8	0.6	0.9	0.1
Dissimilarity index	–	0.47	0.47	0.39	0.6	0.03

Source: Authors' calculations (ARTIS), Soramäki et al. (2006; FedWire).

Note: The value and average value figures for FedWire are converted into euro based on the USD/EUR exchange rate of 1.21730 of March 31, 2004.

tled payments and 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 topology analysis we focus on the giant strongly connected component (GSCC) of ARTIS.² 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 relevant to financial stability

(e.g. small charities and offset accounts of the 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 those reported in Soramäki et al. (2006) for the GSCC of FedWire.

ARTIS processes on average 15,380 transactions per day, with the daily average value totaling EUR 48.5 billion. The average transaction size amounts to EUR 3.2 million. The size of the network is defined by the number of nodes n . On 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).³ The connectivity p of the

² For mathematical definitions of the network indicators, see the Appendix in Schmitz and Pühr (2007) and Zhou (2003). For comparable data on the network of all active accounts, see Schmitz and Pühr (2007). For a description of the Austrian banking system, see OeNB and FMA (2004, pp. 50–55).

³ The average number of nodes in ARTIS active on every day was 209.8 and these were connected by 1,637.5 directed links.

network is captured by the number of actual directed links relative to the number of possible directed links. Connectivity p averages 7.9%.

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, meaning that it takes only slightly more than two links on average to reach any terminating node in the network from any originating node in the network. Hence, the network is compact, with almost all active nodes 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 maximizing 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 farthest away from it. We thus obtain a value of 4.4 links.

How well are the nodes connected to 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 then averaging across nodes.⁴ Averaged also across days, k amounts to 15.6 in the ARTIS system. In other words, when you pick a node in the GSCC on a random day in the sample period, it can be

expected to have 15.6 links originating (or terminating) at it. A much larger number of links originates and terminates at the most active nodes, however. The maximum out-degree averages 76 across days, so that the most active node of 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 than the average degree. The clustering coefficient provides a measure of the average connectivity of the neighbors of all nodes in the GSCC. On average, about 58% of the neighbors of each node are also interlinked. Betweenness centrality measures how many shortest paths through the GSCC pass through the average node. The value of 0.8% is quite low and stems from the central position 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 any two neighboring nodes. If the network looks very similar from the respective node pairs, the dissimilarity index is close to zero. In the GSCC, it amounts to 0.47, which implies that on average the perspectives of the GSCC differ substantially from any two neighboring nodes. Many nodes are linked to each other although not otherwise sharing many network characteristics. We interpret that as further evidence that many of the nodes connect to the largest nodes at the center of the network.

How do these values compare with the results for FedWire? To begin with, we must bear in mind that the FedWire data refer to the first quarter of 2004. Value and volume in FedWire have certainly grown since then. Comparing a

⁴ The out-degree refers to the number of links originating at the node, while the in-degree is based on the number of links terminating at the node. Across the network, the average out-degree and in-degree are equal to m/n , respectively.


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, implying 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). And indeed, the ratio between connectivity p in ARTIS (7.9%) and that in FedWire (0.3%) is 26:1. 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.⁵ The high

clustering coefficients in both networks (on average 53% vs. 58% of the direct neighbors of each node are also linked) corroborate this finding. The average degrees of both networks are very similar too (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, 2005). 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 BOJ-Net. For the Austrian interbank market, Boss et al.

Chart 1

Histogram and Reverted Cumulated Distribution Function (on Logarithmic Scales) of the Degree Distribution in the Monthly Network in ARTIS (GSCC)



Source: OeNB.

⁵ In a small-world network, most nodes can be reached from each other by a small number of hops or steps, although connectivity is low and most nodes are not neighbors.

(2004) 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 network⁶ (degree range 1 to 1,925 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 left-hand panel of chart 1). 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 preserves 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-hand panel of chart 1. Obviously, the cdf is not a straight line and we reject the power law hypothesis for the ARTIS network.

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; the data cover the period from 2000 to 2003). As they are settled through ARTIS, interbank market transactions can be interpreted as a subset of the transactions processed by 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 . This similarity arises because both the interbank

market and the payment system are dominated by large banks. In both markets, many banks cluster around their sectoral central institutions.⁷ However, the clustering coefficient is substantially higher in ARTIS than in the interbank network. Maintaining interbank relationships is costly, so 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% of the total value). These reflect the network structure of real economic activity, which does not necessarily mirror the structure 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 with roughly 650 million transactions from November 16, 2005, to November 16, 2007 (excluding Austrian holidays).⁸ These simulations were calculated with a self-implemented Matlab-based software tool (inspired by the Bank of Finland Payment System Simulator), which was tailored to ARTIS particularities. The tool recalculates each day's transactions by adding incoming payments to and subtracting outgoing payments from the respective accounts of the participants. As the

⁶ We conducted the same exercise for the daily, the quarterly and the semiannual networks with the same results.

⁷ 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% of Austrian banks in terms of the number of credit institutions and for about 50% 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 have direct access to the system based on their own in-house systems. Although IT solutions within sectors are often similar, there is no evidence that operational risk is correlated across individual banks within a sector.

⁸ For more details on the simulations, their motivation and design, see Schmitz and Pühr (2007). The operation of ARTIS was discontinued after November 16, 2007, due to the introduction of TARGET2.

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 system participants' behavioral reactions into account, they must be determined exogenously. First of all, system participants might want to stop submitting payments to the participant experiencing operational problems. A stop-sending rule applies if a transfer account of a central bank in the TARGET system experiences an operational problem, i.e. no further payments are transferred to the stricken transfer account.⁹ Payments to other participants are not affected. In cases of operational problems at other banks, ARTIS operators at the OeNB provided evidence that participants continue to submit payments to the affected participants, 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 ARTIS participants already hold large shares of their eligible assets at the

OeNB. Consequently, we assume that system participants are not increasing collateral for operational incidents with durations 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 the OeNB as the binding liquidity constraint for banks. Third, the simulation algorithm takes into account debit authorization by banks for a number of other participants in ARTIS.¹⁰

The scenarios in Schmitz and Pühr (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.¹¹ This resulted in three scenarios: in the first, the most active transfer account¹² was shocked; in the second, the most active bank account was assumed to experience operational problems; and in the third, 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

⁹ Due to the operating procedures, it actually takes about 40 minutes following the detection of the operational problem at the transfer account until a stop-sending rule is imposed. The implementation of the rule in the simulation algorithm takes this small delay into account.

¹⁰ Participant A may grant participant B a debit authorization according to the Terms and Conditions Governing the OeNB's ARTIS System (Article 9). Debit authorization is defined as the right of participant B to initiate (certain pre-agreed) payments from the account of participant A. Debit authorizations are granted to a small number of participants for prearranged purposes (very frequent recurring standard operations) and cannot be interpreted as a crisis mitigation instrument available on short notice in the case of an operational incident.

¹¹ 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 (based on both the number and the value of payments received and submitted) as well as (4) the monthly network topology.

¹² Transfer accounts are ARTIS accounts held by other ESCB central banks at the 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.

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.¹³ This assumption is extreme but plausible. Shorter outages of participants may lead to payment delays but not to unsettled payments, as shown in Schmitz and Pühr (2007).

The results are graphically represented in the four panels of chart 2. 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 number of simulations that yield x contagious defaults on the y-axis. It is evident that about 27% of all simulations (8,604) do not lead to contagion at all. Another 26% (8,230) yield one contagious default and 16% (4,919) two such defaults. About 29% (5,456) lead to three to five contagious defaults and 17% (4,102) to more than five. The 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% of the mean. This motivates the investigation in subsection 5.1 as to whether the variation of network topology across days can contribute to the explanation of the fluctuations of average contagious defaults per day.

The lower panels in chart 2 show the average contagious defaults per sim-


ulation (in terms of the 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. 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. This threshold value, which can be measured by the number of contagious defaults or by the value of unsettled payments in the system, is somewhat arbitrary and depends on the risk aversion of the supervisory authority. Setting the threshold in terms of the number of contagious defaults at 1 (to capture 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. This figure includes 11 transfer accounts operated by central banks (lower left-hand panel of chart 2). The 28 bank accounts constitute 12% of the average of 230 bank accounts in ARTIS (during the sample period) and represent about 3% of the average of 850 banks in Austria. Defining the threshold in terms of the value of contagious defaults to pinpoint only accounts that cause at least an average value of EUR 48.5 million of unsettled payments (or 0.1% of the average value of transactions settled across days), we find that 24 accounts are systemically relevant

¹³ 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 Pühr (2007), but disregard them in this paper, as they are of little relevance for the interaction between network topology and contagion.

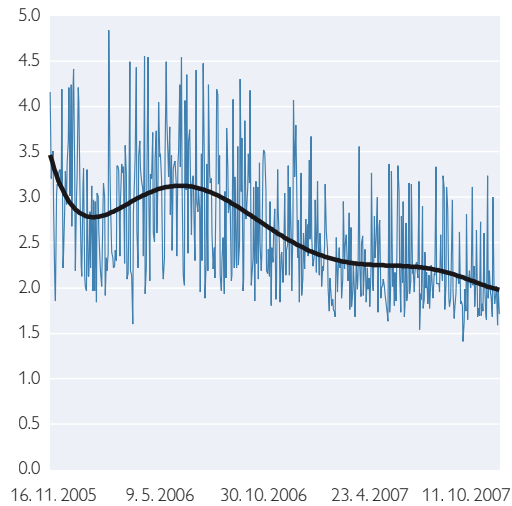
Chart 2

Simulation Results


Number of simulations (in thousands)




Average number of contagious defaults per day across simulations



Average number of contagious defaults per day



Average value of contagious defaults in EUR million



Source: OeNB.

(lower right-hand panel of chart 2). Seven of these are transfer accounts, which leaves 17 bank accounts, which account for about 7% of the average of 230 bank accounts in ARTIS (during the sample period) and for around 2% 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 considerably reduces the liquidity sink effect, the strong contagion impact of transfer accounts is interesting. This indicates that payment concentration risk is more important for the contagion impact than liquidity concentration risk. The fact that TARGET2 operates on a Single Shared Platform without highly contagious transfer ac-

counts might 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.

4.1 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 be able 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 2 (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 did actually 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 discretized 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. Discretizing 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 to describe the observed number of contagion events because 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).¹⁴ 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. For simulations that show at least one contagion event, we conclude that the Weibull distribution is a reasonable choice for describing the probability that the number of observed contagion events C is equal to a positive integer n given by:

$$P\{C=n\} = Wei(n|a,b) - Wei(n-1|a,b) \text{ for all } n \geq 1, \\ \text{where } Wei(\cdot|a,b) \text{ denotes the cumulative distribution function of a Weibull distribution with parameters } a \text{ and } b, \text{ defined by} \\ Wei(x|a,b) = 1 - \exp(-(x/a)^b) \text{ for all } x \geq 0.$$

¹⁴ 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 of statistical tests (DeGroot, 1985).

We thus approximate the distribution of the number of contagious defaults in a simulation given that contagion actually by means of a discretized 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 (subsection 5.1) and at the node level across stricken accounts (subsection 5.2) explains the variation of contagion across days and across stricken accounts.

Selecting 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 indicators available 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 how the flowing good is passed on along the route from one node to another. While a disease can be 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 this 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. Instead, they transfer directly to the ultimate receiver. However, the flow of liquidity does not stop there. Liquidity 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. This 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.

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 (such as betweenness centrality). Third, we want to investigate whether network indicators in general add value to the more traditional mea-

Chart 3

Value and Volume (Network Level) per Day versus Average Number of Contagious Defaults per Day



Source: OeNB.


sure 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 3 we depict the daily value (left-hand panel) and the daily 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 varia-

Chart 4

Selected Network Indicators (Network Level) per Day versus Average Number of Contagious Defaults per Day



Source: OeNB.

tion of value explains 2% and the variation of volume accounts for 8% of the variation of the contagion impact per day.

The explanatory value of the variables value and volume is low. Do network indicators perform any better? In chart 4 we look at the following indicators (unweighted, undirected): degree, average path length, betweenness centrality and dissimilarity index. Similarly to chart 3, the daily number of contagious defaults (in terms of the number of banks with unsettled payments) is depicted on the x-axis and the daily values of the respective network indicator are shown on the y-axis in each panel.

The average path length (15%) and betweenness centrality (13%) have the highest explanatory values. The daily variation in degree accounts for 10% of the variation in contagion and that of the dissimilarity index for only 3%. 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 explana-


tory power of any of the remaining 39 indicators is 15.4% (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, this does not preclude that structural differences across networks might influence a network's 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 of chart 3). Do the different positions of the nodes (that experience the operational shock) in the network account for this variation? In chart 5 we plot the value and volume of payments of the stricken node in each simulation against its contagion effect in terms of the number of contagious defaults (in terms

Chart 5

Value and Volume (Node Level) per Stricken Account versus Number of Contagious Defaults per Simulation



Source: OeNB.

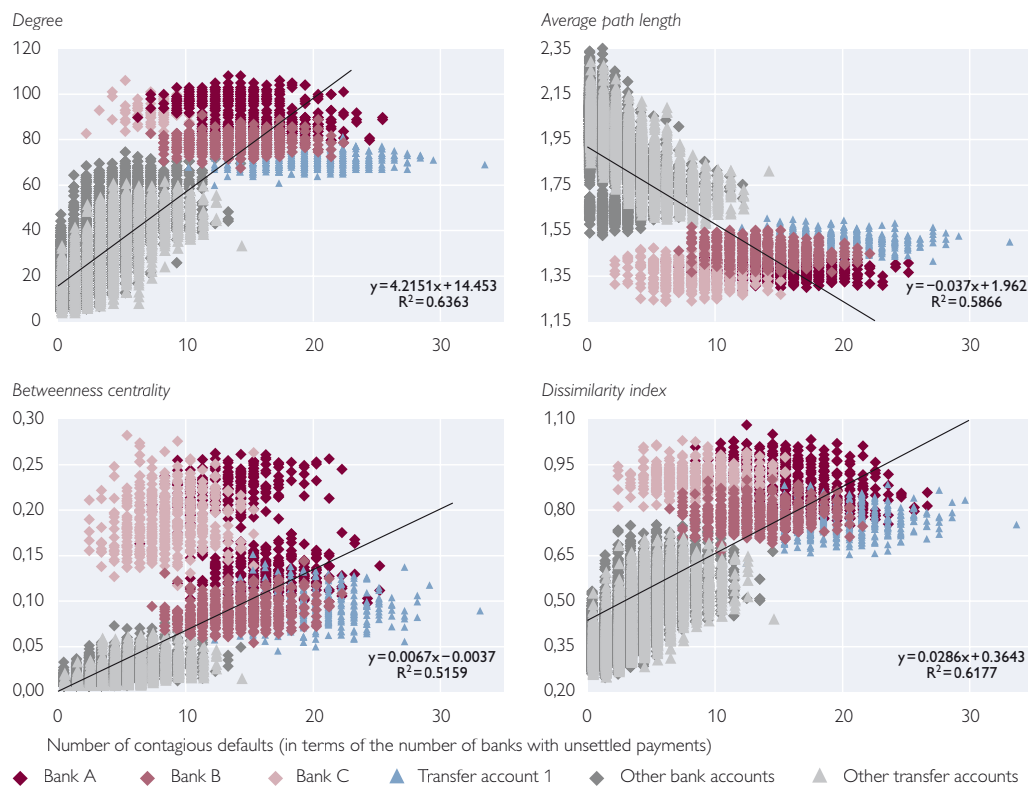
of the number of banks with unsettled payments), i.e. each panel contains 31,311 data points. In addition, the data points of the three most active banks (Bank A, B, and C) and of the most active transfer account (Transfer Account 1) are colored (see the legend of chart 5), while those of all other bank accounts and of all other transfer accounts are dark grey and light, respectively. The variations of value and volume across simulations explain 73% and 68% 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 cause more contagion than estimated by their volumes of transactions) and Banks A and C above the regression line (i.e. they cause less contagion than estimated by their volumes of transactions).

In chart 6 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 the number of contagious defaults (i.e. each panel 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 6

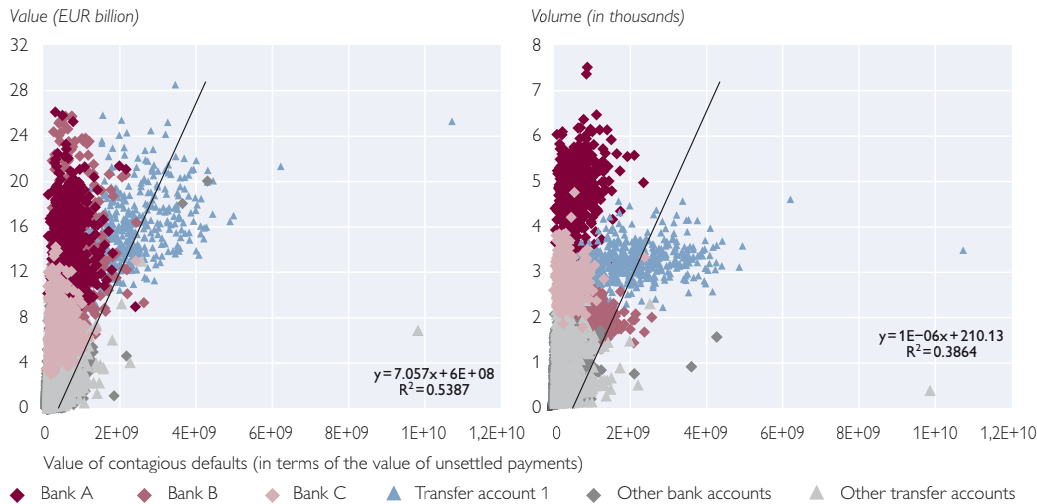
Network Indicators (Node Level) per Stricken Account versus Number of Contagious Defaults per Simulation



Source: OeNB.

Chart 7

Value and Volume (Node Level) per Stricken Account versus Value of Contagious Defaults per Simulation



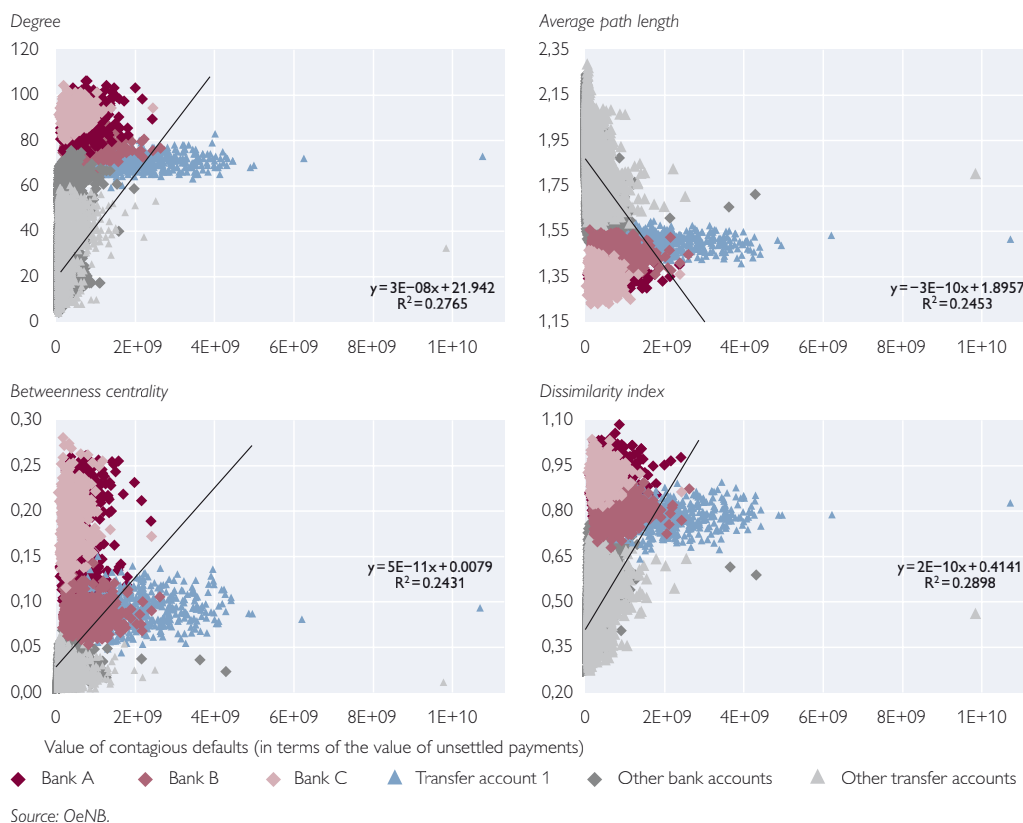
Source: OeNB.

chart 5. The explanatory values of all four network indicators are quite high; the simplest measure degree yields an R^2 of 64%, variations in average path length across simulations account for 59% of the variation of the number of contagious defaults across simulations. The more complex measures betweenness centrality and dissimilarity index yield an R^2 of 52% and 62%, 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% (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 ac-

counts. However, network indicators 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 a data mining problem. The differentiation according to the stricken account confirms the pronounced grouping evident also in chart 5. 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 the left of the regression line. This finding points to structural differences in contagion impact, which are not accounted for by measures of activity or network indicators and warrant further research.

We also investigate the interaction between network topology and network stability for another measure of contagion, namely the value of unsettled payments. Again we start with the analysis of the explanatory value of node size, i.e. of value and volume of

Network Indicators (Node Level) per Stricken Account versus Value of Contagious Defaults per Simulation



payments originating at the node (chart 7). Variations in value explain 54% and in volume 39% of the variation in contagion. Both values are lower than the respective results in chart 5.

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%, average path length 25%, betweenness centrality 24% and dissimilarity index 29%, chart 8) and they are considerably lower than the respective values for the measures of size in chart 7. When 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 7 and 8. Comparing the results for the two measures of contagion, number of banks with unsettled payments (charts 5 and 6) versus value of unsettled payments (charts 7 and 8), 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.

Table 2

Correlations Between Network Indicators (Node Level)

	Volume	Value	Average path length	Degree	Connectivity	Clustering	Betweenness centrality	Dissimilarity index
	%							
Volume	100	89	-77	84	83	-57	89	85
Value		100	-70	76	75	-52	77	78
Average path length			100	-96	-97	62	-79	-85
Degree				100	99	-72	85	95
Connectivity					100	-72	85	93
Clustering						100	-56	-78
Betweenness centrality							100	87
Dissimilarity index								100

Source: OeNB.

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 reveal that various indicators of centrality (average path length, degree, connectivity, betweenness centrality and dissimilarity index) are highly correlated with value and volume.

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, mostly because almost all active nodes are linked to a small number of accounts at the center of the network (the largest banks and the most active transfer accounts). This network structure is quite stable across days. Comparing the ARTIS system with 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 ARTIS with those of the Austrian interbank market reveals that the distance measures are very similar, while the clustering coefficients differ substantially. The similarity arises because the interbank market is likewise dominated by a few large nodes at the center of the network.

We conducted 31,311 simulations based on 63 different scenarios for 497 transaction days from November 16, 2005, to November 16, 2007 (excluding Austrian holidays). Although the scenarios focus only on the banks and transfer accounts represented 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-fifths 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 no more than 28 bank accounts, but almost all transfer accounts operated by central banks, are systemically relevant. If we define systemic relevance as a contagion impact of at least 0.1% of the average value of transactions settled across days, we find that 17 bank accounts and 7 transfer accounts are systemically relevant. In both cases only 7% to 12% of all bank accounts in ARTIS and 2% to 3% of all Austrian banks are systemically relevant. The simulation results suggest that the ARTIS system is remarkably 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 removing transfer accounts in the Single Shared Platform of TARGET2 can improve the system's resilience compared with 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 of our research, network indica-

tors 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 as opposed to the value of unsettled payments. It is questionable at this stage whether network indicators contain much additional information compared with value and volume, which have traditionally been 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 for controlling for other explanatory variables, such as beginning-of-day liquidity at the network and at the node level.

Bibliography

- Albert, R., H. Jeong and A.-L. Barabasi. 1999.** Diameter of the World Wide Web. *Nature* 401. 130–131.
- Albert, R., H. Jeong and A.-L. Barabasi. 2000.** Error and Attack Tolerance of Complex Networks. *Nature* 406. 378–381.
- Albert, R. and 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 and 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 and 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. and S. W. Schmitz (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. <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.
- OeNB and FMA. 2004.** The Austrian Financial Markets. Vienna.
- Schmitz, S. W., C. Pühr. 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. and C. Pühr. 2007.** Risk Concentration, Network Structure and the Assessment of Contagion in the Austrian Large Value System ARTIS. In: Leinonen, H. (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. and G. E. Wood (eds.). 2006.** Institutional Change in the Payments System and Monetary Policy. London: Routledge.
- Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass and 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 and R. J. Glass. 2007.** New Approaches for Payment System Simulation Research. In: Leinonen, H. (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.