



# From agent-based modeling to actor-based reactive systems in the analysis of financial networks

Silvia Crafa<sup>1</sup> 

Received: 27 July 2020 / Accepted: 9 February 2021 / Published online: 20 March 2021  
© The Author(s) 2021

## Abstract

We present a new framework for the analysis of financial networks, called Actor-based Reactive Systems (ARS), that pushes further the Agent-Based approach (ABM) by resorting to ideas coming from the study of distributed systems in computer science. Two distinctive features, namely a fundamentally different management of time and a fully decentralized control logic, have a profound impact in terms of expressiveness of analysis, flexibility of modeling, and efficiency of experimentation. To illustrate the feasibility of the framework, we develop a realistic case study by analyzing the systemic risk of a model of the European banking network with a nontrivial contagion procedure, that combines an initial asset shock with the negative feedback loop triggered by asset fire sales. We show that, compared to ABMs, ARSs bring about finer-grained analyses, with a greater degree of heterogeneity and adaptivity of economic agents. Moreover, the very low computational cost and the detailed account of the system's execution support the design and the development of very flexible stress tests to rapidly experiment with many hypothetical scenarios in a test-oriented style.

**Keywords** Systemic risk · Agent-based models · Distributed systems · Computational models · Reactive systems

## 1 Introduction

The behavior of financial systems is typically studied in terms of graph-based models, where nodes represent financial institutions and links represent financial dependencies, like interbank exposures. A large body of literature applies mathematical methods to the analysis of global properties of the financial networks, and in particular systemic risk, that is the fragility of the system when exposed to the contagious effects of distressed financial institutions (Cont et al. 2013; Cont and Minca 2016; Capponi

---

✉ Silvia Crafa  
silvia.crafa@unipd.it

<sup>1</sup> Dipartimento di Matematica - Università di Padova, via Trieste 63, 35121 Padova, Italy

and Bo 2015; Capponi 2016; Battiston et al. 2012, 2016; Neveu 2018; Dai Pra et al. 2009). While graph modeling enables rich mathematical descriptions and sophisticated analysis techniques, modeling—thus thinking—in terms of graphs, naturally shifts the view towards a *global view* of the system and to its *static structure* rather than its dynamics. Accordingly, the literature associates the systemic risk to a number of structural properties, i.e., network-based measures of connectivity and risk, such as counterparty susceptibility, local network fragility, feedback centrality, DebtRank. On the other hand, the economic crisis in 2008 called attention to the finer-grained study of the interactions of economic agents, so to better account for the global system's dynamics emerging from the local components' behavior. This is the approach followed by Agent-Based Models (ABMs) of economies (Tesfatsion and Judd 2006; Farmer and Foley 2009; Trichet 2010; Seppelcher 2012; Teply and Klinger 2019). These models usually represent static network structures where nodes, called agents, have a dynamic behavior expressed in terms of rules that can be executed by a system simulator, i.e., a computer program. Compared to graph models, agent-based models have a more dynamic nature, but they still encode a rather global view of the system. Indeed, even if heterogeneous behavior rules can be assigned to different (sets of) agents, the system is generally governed by a *centralized logic*, where the agents' behaviors are mathematical functions that depend (also) on global variables, like the current market state or the global timing. In particular, the time flows synchronously within the system, whose evolution is given in terms of a sequence of *global and synchronous iterations*. At the beginning of each simulation iteration, every agent re-executes its local behavior function using the updated values of the variables that have been modified at the end of the previous iteration. Even if each iteration is assumed to last for an unspecified period, a modeling framework with this level of synchronicity and centralization cannot account for situations where different financial institutions react to the same event (e.g., a macroeconomic shock or a bank default) at very different times (e.g., a bank may quickly absorb the shock, another one might need time to liquidate part of its assets, and different countries might impose different procedures for writing off bad debts).

Therefore, Crafa and Varacca (2018) proposed a novel conceptual framework that pushes further the ABMs with (i) a fundamentally different management of time, and (ii) a fully decentralized control logic. The new framework is borrowed from the field of computer science and ICT, in that it models financial systems as distributed and reactive systems, that is a well known class of *asynchronous* communicating systems, characterized by a remarkable level of *decentralized control* and *coordinated interactions* between agents located on different sites. The key point is that financial entities operate in parallel and dynamically update their state and behavior as a reaction to events issued by other entities. Economic agents are only partially aware of the current global state of the system; they asynchronously receive messages that notify that an event has occurred (e.g., a market shock or a debtor communicating its default), so that they can react to that event by modifying their behavior according to their internal policy and locally stored information (*strong locality and isolation* principles). We call this approach Actor-based Reactive System (ARS), both to remark its distinctive features and to allude to the specific model of information systems on which it is based, that is the so-called Actor Model (Agha 1986).

In this article we show the feasibility of the ARS framework by putting it at work on a realistic scenario: the systemic risk analysis of the European banking network. We remark that our aim here is not to provide financial insights on the Eurosystem stability, nor to precisely estimate the systemic risk of specific local banking sectors. This would require a more experienced parameter calibration, because the system behavior is significantly influenced by the model parameters, as it is usual in the literature about systemic risk analysis. Instead, we call attention to how the ARS, compared to ABM, brings about finer-grained analyses and more flexible modeling, with a greater degree of heterogeneity and adaptivity of economic agents.

In particular, we consider an European banking network closely related to that studied by Tepy and Klinger (2019) using an Agent-Based Model. It is a network of 286 banks, in 9 European countries, with structured balance sheets of different sizes. Realistic interbank exposures are reconstructed using available data from the European Central Bank's statistics website and the Bank for International Settlements' consolidated banking statistics. We consider a nontrivial contagion procedure, that combines an initial shock on external assets with the negative feedback loop triggered by asset fire sales (as Tepy and Klinger (2019) or Battiston et al. (2016)). However, the centralized logic governing the contagion cascade in agent-based and graph-based models is rephrased here in terms of a fully parallelized and decentralized logic, where actors asynchronously react to messages that notify either an interbank contagion or a funding liquidity shock.

We then use the FinMarkSim tool developed by Crafa and Varacca (2018) to directly run the ARS model under various conditions and test the role of different parameters on the systemic risk of the European network. We study the influence of the size of the initial asset shock and the market illiquidity factor, the systemic influence of specific countries and bank groups, and the influence of the topology of the financial network. We show that, by studying the system's evolution as an Actor-based Reactive System instead of a sequence of synchronous iterations, it is possible not only to assess the final global state of the system, but also to precisely inspect the complete and fine-grained dynamics of each agent, observing the *causal chain of events-and-reactions* that determined the agent's behavior within the system execution. As a consequence, this approach allows to precisely assess and disentangle the roles of two different contagion channels, i.e., the interbank contagion and the illiquidity channel, within different countries or in specific banks. Moreover, since each complete system's run lasts up to 3–6 seconds (most tests handle the behaviors of more than 1400 agents in 2–3 seconds), the FinMarkSim tool can be effectively used to design and conduct fine-grained stress tests with hypothetical scenarios in a test-oriented style.

The paper is organized as follows. Section 2 discusses how to model a financial system as an ARS, providing an accessible explanation of the model's features borrowed from computer science. Sections 3, 4 describe the structure and the dynamics of the ARS corresponding to the European banking network that is studied in depth in this article. Section 5 reports on the stress tests, and Sect. 6 concludes the paper with final remarks.

## 2 A financial system as an actor-based reactive system

In this section we overview the general architectural pattern proposed by Crafa and Varacca (2018) to represent a financial system as an ARS, focusing on the differences with the ABM. This general pattern will be then instantiated in the next sections to model the European interbanking network.

### 2.1 Economic actors

As in the agent-based approach, the core of actor-based reactive systems is a set of individual components, called actors or processes, that run autonomously in parallel and explicitly interact according to their local behavior/reaction rules. We consider two kinds of actors: financial *institutions* (e.g., banks, finance companies, brokerage companies, development agencies) and financial *contracts*, which generally refer to any financial agreement (e.g., loans, credits, bonds, derivatives like credit default swaps) and any kind of assets (e.g., sovereigns, households, physical assets).

In financial models contracts are usually described by local variables of the agent that represents the bank that holds the contract. Instead, considering contracts at the same abstraction level of institutions (i.e., they are both economic actors) is a key choice of the architectural structure of Crafa and Varacca (2018), that leads to a better systemic account of the financial relationships. Indeed, by decoupling contracts from institutions and letting them run autonomously and in parallel with other agents, we have that more agents' interactions are possible and thus additional system's behaviors can be captured and more deeply analyzed. For instance, if loans run in parallel with both the debtors and the creditors, the system can account for the parallel progressing of the contagion through different channels and the different recovery strategies that can be adopted by the two sides of the financial contracts.

### 2.2 Actor behavior and interactions

The main difference between ABMs and reactive systems is in their dynamics. While an agent-based model is typically iterated over time steps, with aspects of the agents updated at each time step, actors autonomously run in parallel, occasionally interacting by exchanging messages. The sending of a message can be thought of as the communication that an event has occurred (e.g., a market shock or the default of a debtor), so that the receiver of the message can *react* to the received message by modifying its behavior according to its internal policy and locally stored information. The actor behavior is then defined as a set of messages that the actor can receive (*interface*) and the corresponding reactions performed by the actor upon receiving those messages. Reactions involve the update of internal information, some local computation and the sending of other messages to other actors.

## 2.3 System's dynamics

Actors' communication is generally asynchronous and point-to-point, but synchronous messages, that block the sender until the receiver replies, and broadcast are occasionally useful, e.g., to signal a global fire sale. Notice that asynchronicity seems to be particularly well-suited to financial interactions, where counterparties tend to take their time to react to requests, e.g., a request of payment can be fulfilled after some delay, either because the debtor needs time to collect liquidity, or because of strategic or market reasons. Exposing such a delay in message responses allows the model to capture complex scenarios. As an example, consider the case of a bank  $b$  that receives a message from a debtor that notifies its default, and reacts to the credit loss by liquidating part of its assets. While  $b$  is collecting liquidity, it might receive a message that notifies a global macroeconomic shock that decreases the value of assets. In the system's run the message that notifies the asset shock can be received by  $b$  either before or after the message that notifies the completion of the asset liquidation procedure, thus possibly resulting in either the default or the saving of the bank  $b$ . A further example is a scenario where financial institutions' portfolios also contain Credit Default Swaps derivatives (CDS) or other forms of insurances against credit risk. Upon the default of the CDS's obligor, the buyer of the CDS expects to receive the corresponding notional; however, in a situation of general distress where the buyer needs to collect liquidity, a significant delay in having at its disposal the CDS notional may increase the stress and possibly lead to the buyer's failure (Crafa and Varacca 2018).

In other terms, each system's run schedules a *nondeterministic interleaving* to decide in which order parallel messages directed to the same agent are received by that agent. Then, different system's runs might exhibit different interleavings of parallel actions, thus different system evolutions, while *respecting the causality chains of events*. Unless explicitly programmed, by tuning system parameters, all the message interleavings are equally possible. Therefore, by running the system several times, different system's behaviors can be observed, and the causal chains of events that led to those behaviors can be precisely traced and analyzed by resorting to the verification tools provided by the theory of distributed systems.

We will see in the next sections that an inspection of the ARS execution reveals the precise causal chain of events that led from a global market shock to a specific bank's default or distress. Thus, precisely tracing the chains of events-and-reactions allows one to investigate specific interbank contagion channels between different countries or between different banks size-groups. More generally, the ability to model and reason about nondeterministic behavior is particularly relevant in all those situations where the path (i.e., the sequence of events) that leads from an initial distress to other bank defaults might be relevant to devise a recovery strategy.

## 2.4 System's execution

Instead of mathematically simulating a (agent-based) model, an ARS is a *computational model* that can be directly encoded and run as a computer program, allowing a more direct and finer-grained control of the behavior of the system's components.

**Table 1** Banks data

	Market concentration	Large banks	Medium banks	Small banks	Total banks	Total assets
Austria	Low	4	16	24	44	840
Belgium	Medium	6	6	10	22	989
Germany	Low	4	16	24	44	7773
Spain	Medium	6	9	18	33	2638
France	Medium	6	6	10	22	8477
Ireland	Medium	6	6	10	22	560
Italy	Low	6	9	18	33	3666
Netherlands	High	6	9	18	33	2315
Portugal	High	6	9	18	33	391
Total					286	

The market concentration index and the number of banks are taken from Teply and Klinger (2019). Total country assets (in EUR billions) are taken from the ECB document about aggregated balance sheet of euro area credit institutions (ECB 2018)

**Table 2** Market shares (from Teply and Klinger 2019)

Market concentration	Herfindhal index	Small banks share (%)	Medium banks share (%)	Large banks share(%)
Low	[0...500)	45	15	40
Medium	[500...1100)	15	25	60
High	[1100...10000]	10	20	70

Crafa and Varacca (2018) put forward an architectural pattern (patent pending) to be used to appropriately encode financial networks as distributed systems based on the Actor Model (Agha 1986). Its direct software implementation, called FinMarkSim, can be used to build the financial system directly in a programming language and execute it. It thus provides for a very flexible and interactive stress test framework, that supports the analysis of a financial network in a test-oriented style, which is hardly possible with other more complex mathematical simulation methods.

### 3 The European banking network

We now illustrate a financial ARS at work on a real-world case study. To better compare our approach with ABMs, we consider the European banking network studied by Teply and Klinger (2019), using available data from the European Central Bank's statistics website and the Bank for International Settlements' consolidated banking statistics. Our model is not the precise replica of that network since a number of model parameters were not available; moreover we consider different graph topologies to investigate the role of different patterns of node interconnectedness.

**Table 3** Balance sheets of large, medium-sized and small banks (Teply and Klinger 2019)

	Assets (%)			Liabilities (%)		
	Sovereign debt	Interbank	External	Interbank	External	Equity
Large	15	25	60	25	70	5
Medium	20	10	70	10	80	10
Small	30	5	65	5	80	15

### 3.1 The bank modeling

We consider a system of nine European countries, whose banking systems are approximated using their home country parameters, according to structural financial indicators available from ECB and following the modeling method of Teply and Klinger (2019), which is in accordance with the patterns in the balance sheets of real-world banks. Table 1 shows the list of countries, the number of banks, classified in three categories: small-sized, medium-sized and large banks, the outstanding amount of total country assets (in EUR billions) and the market concentration given in terms of the country's Herfindhal index. For each country, the total country assets are split into the assets of each size group according to the prototype market shares given in Table 2. Then, the assets are further divided into assets of each bank depending on the bank number in each size group. Finally, the individual banks' balance sheets are calculated according to the shares given in Table 3, where the total assets for a bank are the sovereign debt, i.e., the sum of a bank's exposure to the local sovereigns, the interbank assets, i.e., the sum of the exposures toward the other banks in the system, and the external assets, i.e., the sum of the bank's exposure outside the banking network, such as households, foreign sovereigns, and non-European financial institutions. The total liabilities are given by the interbank liabilities, i.e., sum of the debts to other banks in the system, external liabilities, i.e., deposits and exposures related to outside the network, and the equity, i.e., the capital reserve representing the security buffer for covering potential losses. Figure 1 shows the aggregate balance sheet structure for different countries and different size groups. For instance, observe that large banks have high interbank assets and low capital buffer, whereas small banks have low interbank assets, and higher equity.

### 3.2 The interbank network

Reconstructing the interbank lending network is much more difficult since full data on mutual exposures of banks, both local and foreign, are not available. This is a well-known general challenge, that is acknowledged by the large body of literature about network reconstruction (Anand et al. 2018; BCBS 2015; Ramadiah et al. 2020; Cimini et al. 2019; Gandy and Veraart 2019). On the other hand, our aim is not to precisely estimate the systemic risk of the Eurosystem, but to show the feasibility of the ARS framework, which is independent of the underlying network topology. Therefore, we resort to a simple but nontrivial reconstruction procedure that keeps the properties of

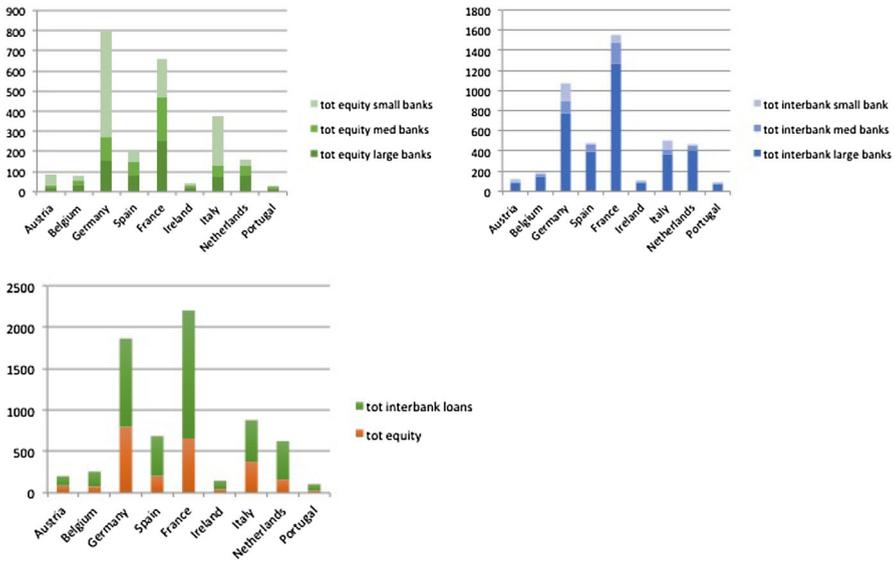


Fig. 1 Aggregate balance sheet structure for banks in the network

Table 4 Probability of connecting two national banks according to their size group

	Large	Medium	Small
Large	0.25	0.10	0.05
Medium	0.21	0.08	0.03
Small	0.17	0.06	0.02

These values ensure that large banks are much more interconnected than small banks, and that small banks act more as lenders while large banks act more as borrowers

realistic financial networks, that is the nets we build tend to be disassortative (i.e., low degree nodes are more likely connected with high degree nodes), sparse, clustered at the level of different countries, and with a few core financial institutions tightly interlinked (Anand et al. 2018; Ramadiah et al. 2020).

We proceed as follows: we first define the topology of the network, i.e., the edges of the graph, and then we define the amount of the exposures using data from the banks’ balance sheets above. On the basis of the observation that the balance sheets structure shows that large banks are much more exposed towards other banks than smaller banks, we let small and medium-sized banks to be exposed just to the national banks, while large banks to be exposed also to the foreign institutions. Therefore, we first define the structure of national interbank nets, and then we connect the country graphs through the nodes corresponding to their large banks.

Local interbank nets are built according to real-word patterns of interbank markets: large banks are much more interconnected than small banks, and small banks act more as lenders while large banks more as borrowers (De Masi et al. (2006); Iori et al. (2008); Soramaki et al. (2007)). Accordingly, we build country networks as random

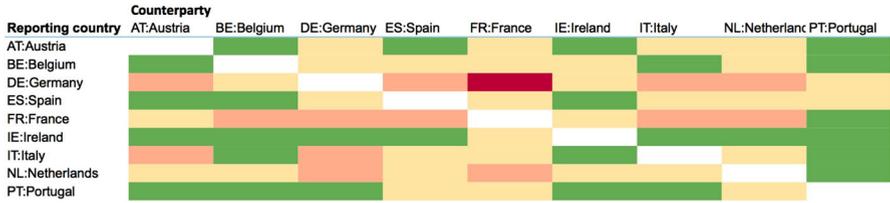


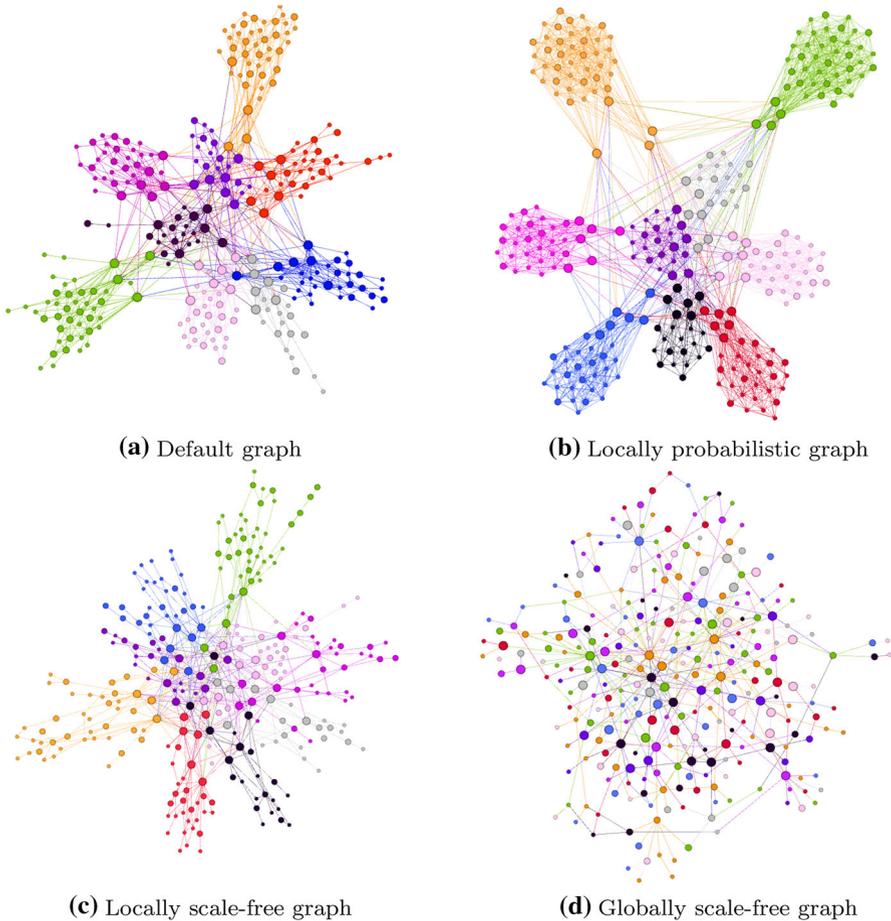
Fig. 2 Foreign exposures heatmap (from Tepyly and Klinger (2019)). We let the green level of exposure correspond to probability 0.06, the yellow level to probability 0.08, the orange level to probability 0.15 and the red level to probability 0.25

graphs where the probability of connecting two (national) banks depends on the size of the two banks. The probabilities, given in Table 4, are input parameters of the model, and are defined so that the probability of having an out-going edge from a small bank (i.e., a credit) is higher than of an in-going edge (i.e., a debt), whereas it is the opposite for large banks.

The global graph is then obtained by connecting the individual country graphs through the nodes corresponding to their large banks. More precisely, we let each pair of large banks of different countries  $c_1$  and  $c_2$  be connected with a probability that is proportional to the level of exposure of  $c_1$  toward  $c_2$ . The foreign exposure levels, and the corresponding probabilities, are shown in the heatmap in Fig. 2. They are built from data considered in the model of Tepyly and Klinger (2019), that are grounded on the data reported to the Bank of International Settlements. The resulting European graph is show in Fig. 3a. Moreover, in order to investigate how the role of interconnectedness in systemic risk can be observed through the ARS framework, we consider three alternative graph topologies. In the *locally probabilistic graph* (Fig. 3b) each country graph is generated by connecting two local banks with a fixed probability  $P = 0.2$ , thus independent of the bank size. The global graph is instead obtained by connecting the large country banks as above. In the *locally scale-free graph* (Fig. 3c) each country network is generated as a scale-free graph, that is a graph where the degree distribution follows a power law: few nodes (called hubs) have high degree while most of the nodes have few connections with other nodes. We generated the graph so that hubs generally correspond to large banks in the country. Then, the global graph is obtained by connecting the large country banks as above. Finally, as a sort of case-control, we also consider a *globally scale-free graph* (Fig. 3d) where there are no clusters corresponding to individual countries. As before, we let larger banks correspond to nodes that have higher degree of interconnectedness.

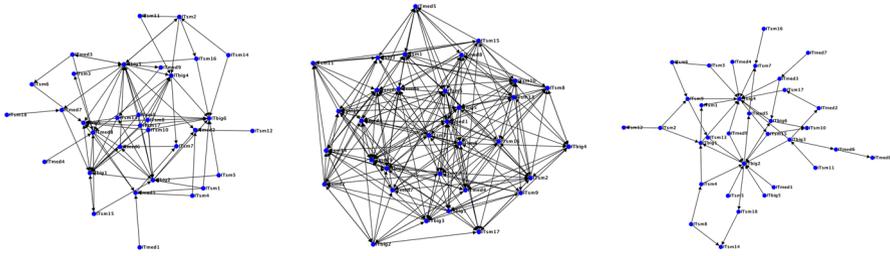
As an example, Fig. 4 displays the Italian banking network that has been isolated from the different graph models.

In these definitions of network we still did not assign any value to the graph’s edges: we just connected the nodes with a number of edges that is proportional to the level of exposures between the European banks, but we need to define the value of the loan represented by each edge. In the balance sheets considered above the interbank assets and liabilities have the same volume (see Table 3); however in banks of different size-groups the loans correspond to different shares; therefore it is not straightforward to choose how to balance each bank’s credits and debts. We adopt the following method:

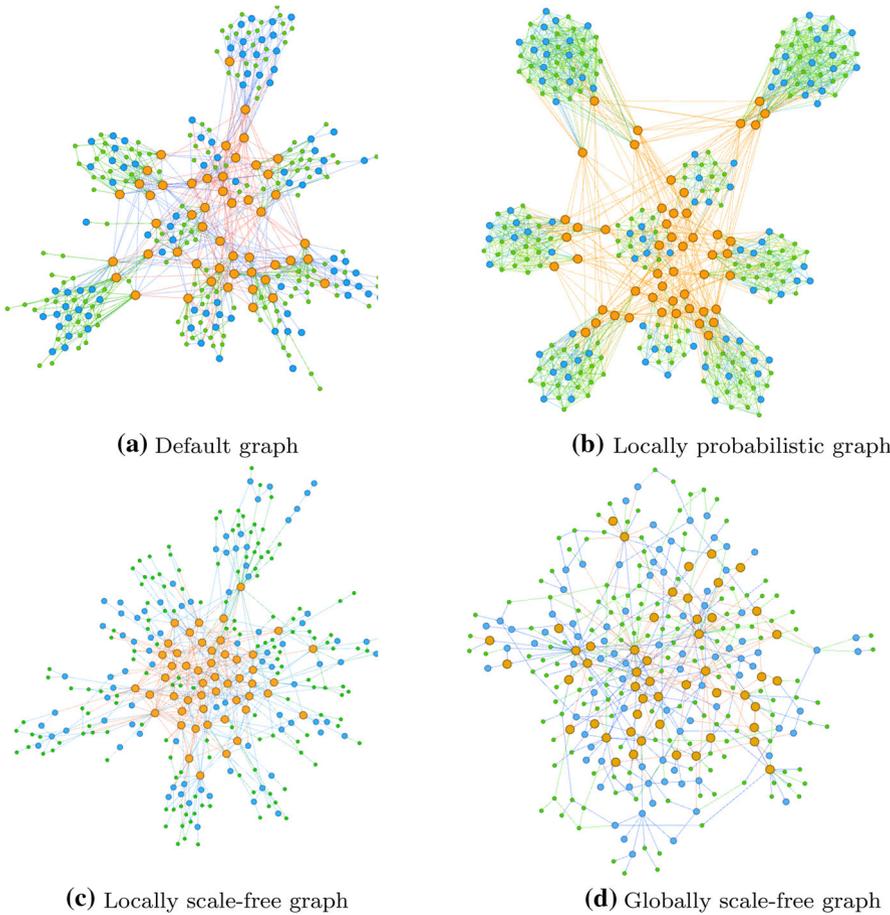


**Fig. 3** Alternative graph topologies. All the graphs have 286 nodes corresponding to the banks in the system. The node size corresponds to the size group of the bank (small, medium-sized, large) and the color corresponds to the country. Different graph topologies have different number of edges: the default graph has 862 edges, the locally probabilistic has 2067 edges, the locally scale-free has 587 edges, the globally scale-free has 434 edges. The edge's color corresponds to the color of the source node, that is the creditor's country. Green nodes correspond to Austria, black to Belgium, orange to Germany, violet to France, red to Spain, blue to Italy, gray to Ireland, pink to The Netherlands and magenta to Portugal

for each edge  $e$  in the graph, let  $b_1$  be the source node and  $b_2$  be the target node. Then,  $e$  represents a loan where  $b_1$  is the creditor bank and  $b_2$  is the debtor bank. Then, we let the value of the loan  $e$  be the minimum between the value of  $b_1$ 's interbank assets divided by  $b_1$ 's out-degree (i.e., the number of  $b_1$ 's credits) and the value of  $b_2$ 's interbank liabilities divided by  $b_2$ 's in-degree (i.e., the number of  $b_2$ 's debts). This approach ensures fair values for edges that connect banks of different size, that have different shares corresponding to interbank assets and liabilities. Finally, we use total assets to calibrate the systems, that is we define the value of the banks' aggregated assets so to restore the balance between the total amount of the bank's assets and the



**Fig. 4** Italian banking network in different graph topologies: to the left the default graph, in the middle the locally probabilistic graph and to the right the locally scale-free graph



**Fig. 5** Different graph topologies with nodes partitioned per size groups and edges partitioned per loan sizes (green for small, blue for medium, orange for large)

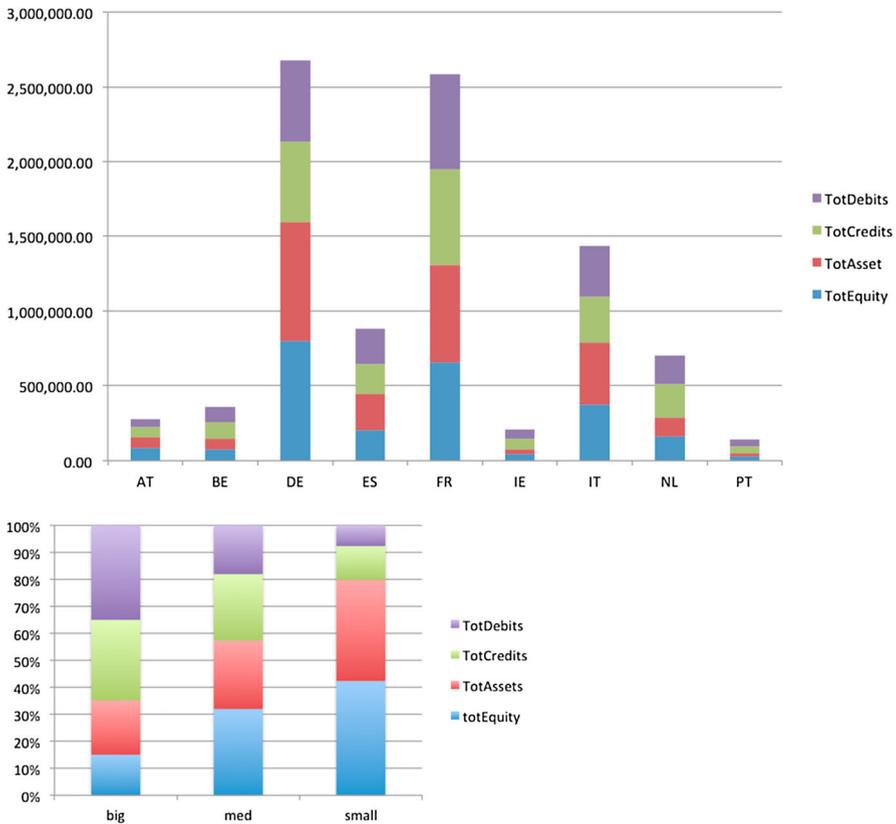


Fig. 6 Summarized balance sheets structure for different countries. Values in EUR millions

bank’s liabilities. Figure 5 shows the networks partitioned so to highlight both the size of the banks and the size of the loans, while Fig. 6 overviews the summarized balance sheets structure for different countries. Observe that credits and debts are well balanced, with small-sized banks acting more as lenders, i.e., their credits are much bigger than debts, while it is the opposite for large banks.

#### 4 The European actor-based reactive system

In the previous section we described the structure of the European interbank network; in this section we focus on the dynamics of the system. First of all we show how to convert the given network into an Actor-based Reactive System based on the architectural pattern described in Sect. 2 and implemented in the FinMarkSim software tool.

The system spawns an Institution actor for each of the 286 banks, and a Loan actor for each edge in the interbanking graph. The internal information of each Institution actor contains the name of the corresponding bank, the value of its capital reserve computed as described in Sect. 3.1, and a list of references to the Loan actors corresponding

to the loans for which the bank is either creditor or debtor, i.e., corresponding to the graph's edges for which the bank is either the source or the target node. Each Loan actor contains the value of the loan, a reference to the Institution actor corresponding to its creditor, and a reference to the Institution actor corresponding to its debtor.

In its generality the architectural pattern of FinMarkSim assumes that a bank may have a list of different assets, each one running independently as a distinct actor of type Asset. This allows one to manage different kinds of assets with different procedures, possibly evolving in parallel. For instance it would be possible to study a scenario where a macroeconomic shock affects the value of fixed assets like buildings, while the assets related to government bonds may be subject to a change in the exchange rates. Moreover, by being independent and running in parallel with the Institution actors, in the meanwhile banks can continue to carry out financial transactions. In this case study we simply consider the aggregated assets of each bank, and we let the system spawn a single Asset actor for each bank. Then, in order to allow communication, the Institution actors also keep a reference to their corresponding Asset actor, and vice versa each Asset actor keeps a reference to the actor corresponding to its owner bank. Summing up, the European ARS considered has more than 1400 economic agents running in parallel: 286 Institutions, 286 Assets and 862 Loans. The behavior of these agents is detailed below; however, we anticipate that the FinMarkSim software completely executes a run of this system in 2–6 seconds, thanks to the scalability properties of the Erlang programming language, used (between others) in computer science to effectively deal with large distributed systems.

#### 4.1 The Actor-based contagion procedure

The aim of the analysis is to evaluate the systemic risk of various interbanking connection patterns. The system's dynamic is then defined in terms of an initial adverse market shock that propagates to the whole system together with a negative feedback loop triggered by asset fire sales (as in Teply and Klinger Teply and Klinger (2019) or Battiston et al. Battiston et al. (2016)). More precisely, the banks are initially hit by the shock, and correspondingly deduce a share of total assets from their balance sheet. As they lose a certain part of their assets, banks write off an equal value of liabilities to balance the assets and liabilities. Banks first attempt to absorb the shock using their equity; however, if the capital buffers are not large enough, they default and the stress propagates through the network until it is completely dissolved, or all the banks in the system are defaulted. Moreover, we assume that banks in default liquidate all their assets in a situation of market illiquidity, where it is not possible for banks to sell the assets for the prices on the bank's balance sheet. The resulting fire sale of assets causes an additional global shock that impacts on the value of the external assets of all the banks in the system. Furthermore, as the failing bank liquidates all of its assets, it may withdraw a certain part of short-term credit claims on other banks. For this reason, the debtors of the failing bank may receive an additional funding liquidity shock and is forced to fire sale an amount of external assets equal to the value of this funding shock.

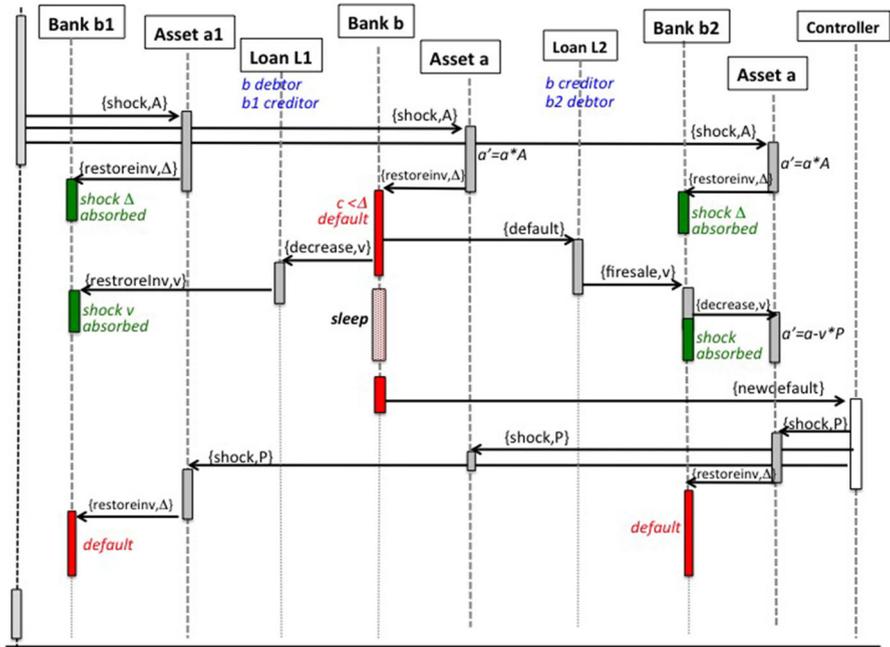


Fig. 7 Sequence Diagram of the execution of a simplified system. Vertical lines denote logical threads: one for each agent. Horizontal arrows denote message sending, and solid boxes denote local computation executed by the actor when triggered by a message reception. In the example bank *b* defaults because of the initial shock, while both banks *b*<sub>1</sub> and *b*<sub>2</sub> absorb the initial shock and the interbank contagion, but they later default because of the price shock due to the default of *b*

Compared to the corresponding implementation in the ABMs, this contagion behavior is rephrased here in terms of reactive actors, with a higher degree of decentralization and asynchronicity, where actors react only upon receiving new messages. On the other hand, when a defaulting bank liquidates its assets, we have that a *local action* (i.e., the bank default) *triggers a global effect* impacting on the assets values of all the banks in the system. Therefore, the system requires a form of global coordination. We implement it by including in the system an additional Controller actor that runs in parallel keeping the references to all the actors corresponding to the non-defaulted banks. Then, any local event, such as the default of a bank, can trigger a global effect by sending a message to the Controller which then broadcasts a suitable notification to all running actors. In this way a centralized coordination logic is implemented asynchronously and in compliance with the principle of locality of information, distinctive of ARS.

#### 4.1.1 Detailed contagion procedure

Figure 7 illustrates the detailed contagion procedure on a simple system that contains only three banks, with assets, and two loans. At the beginning a message {shock, A}

is sent to all *Asset* actors to notify the initial shock.  $A \in [0, 1]$  is a parameter of the model and corresponds to the *initial stress rate*.

Let  $a_j$  be the initial value of the assets of bank  $j$ . Upon reception of the shock message, the asset updates its value to  $a'_j = a_j \cdot A$  and notifies the owner bank  $j$  that a gap of size  $\Delta_j = a_j - a'_j = a_j(1 - A)$  has been created between its assets and liabilities. The bank  $j$  then reacts to the incoming notification (i.e., the message  $\{\text{restoreinv}, \Delta_j\}$ ) by starting a local procedure *Restore\_Invariant*( $\Delta_j$ ) to restore the invariant of its balance sheet due the asset shock of  $\Delta_j$ . The procedure *Restore\_Invariant*( $\Delta_j$ ) tries to locally absorb the shock, if possible, otherwise manages the default of the bank  $j$  and the consequent contagion. It is defined as follows:

- if the capital reserve of the bank is enough, i.e., if  $c_j \geq \Delta_j$ , the bank's equity is reduced to  $c'_j = c_j - \Delta_j$  and the shock does not further propagate (e.g., see banks  $b_1$  and  $b_2$  in Fig. 7);
  - if the equity of the bank is not enough, i.e., if  $c_j < \Delta_j$ , the bank  $j$  defaults (e.g., see bank  $b$  in Fig. 7). As a consequence, the capital reserve is set to zero, (1) the *residual shock*  $\Delta_j - c_j$  flows to the interbank liabilities, (2) debtors and creditors of bank  $j$  are notified for a possible contagion, and (3) the bank  $j$  liquidates all its assets before being removed from the network. Accordingly, the following sequence of actions is executed by the actor corresponding to bank  $j$ .
1. Let  $\ell_j$  be the total amount of bank  $j$ 's interbank liabilities, i.e., its debts. If  $\ell_j \geq \Delta_j - c_j$  it means that the initial shock is absorbed by the bank's capital and interbank liabilities. Otherwise the residual loss of  $(\Delta_j - c_j) - \ell_j$  overflows to external liabilities and must be covered by depositors.
  2. A contagion message is sent to all the *Loan* actors corresponding to  $j$ 's credits and debts, i.e., the actors in the list of loans stored in the bank/actor  $j$ . In order to trace the contagion's causal chain, this contagion message is enriched with a tag recording the source, i.e., the bank  $j$ , and the interbank channel.
  3. Finally, the bank  $j$  liquidates all its remaining assets and sends a  $\{\text{newdefault}\}$  message to the *Controller* actor to trigger the file sale that negatively impacts on the global asset prices.

While the bank  $j$  is executing this default procedure, any other actor in the system is running in parallel according to its internal behavior rules. In particular, the rules governing the reactions to the messages coming from the defaulting bank  $j$  work as follows:

- a *Loan* actor corresponding to a credit of bank  $j$  (e.g., see loan  $L_2$  in Fig. 7) eventually receives the contagion message from  $j$  that notifies the default of the creditor. Then, the loan resets its value  $v$  to zero and forwards the contagion message  $\{\text{firesale}, v\}$  to the actor corresponding to the loan's debtor bank.
- Similarly, a *Loan* actor corresponding to a debt of bank  $j$  (e.g., see loan  $L_1$  in Fig. 7) eventually receives the contagion message from  $j$  that notifies the default of the debtor. Then, the loan decreases its value by the amount of the residual loss that were not covered by the capital buffer, and forwards the contagion message to the actor corresponding to the loan's creditor bank.

- An Institution actor corresponding to the debtor of defaulted bank (e.g., see bank  $b_2$  in Fig. 7) eventually receives the contagion message  $\{\text{firesale}, v\}$  from its defaulted loan. As we said above, we assumed that as the failing bank liquidates its assets, it may withdraw a certain part of short-term credit claims. Therefore, the debtor offers on the market for sale the amount  $v$  of assets to keep the balance, but, because of the fire sale, it receives for it just the amount  $v \cdot P$  which is the price of the assets under the current market valuation. We let  $P \in [0, 1]$  be a model parameter that represents the *price shock*, proportional to the market illiquidity factor, that is the speed at which the asset prices decrease. Finally, the bank tries to absorb the funding liquidity shock and restore the balance between assets and liabilities by invoking the *Restore\_Invariant*( $v \cdot (1 - P)$ ) procedure. Notice that this implies that it can either absorb the interbank contagion or default.
- An Institution actor corresponding to the creditor of defaulted bank (e.g., see bank  $b_1$  in Fig. 7) receives the contagion message from its defaulted loan. This means that one of its debtors defaulted and the corresponding debt has been already decreased of the amount  $v$ . Thus, the creditor actor simply restores the invariant in its balance sheet by invoking the *Restore\_Invariant*( $v$ ) procedure. Notice that this implies that it can either absorb the interbank contagion or default.
- The Controller actor eventually receives a *newdefault* message from the newly defaulted bank that is liquidating its assets. It then reacts by triggering the global firesale in the asset market, that is it sends a  $\{\text{shock}, P\}$  message to all (non-defaulted) assets to notify the price shock along the illiquidity contagion channel. For simplicity we use the same model parameter  $P$  to tune the price shock occurring both in the case of funding liquidity shock and in the global fire sale. A more refined modeling would let the Controller trigger a global shock whose size is proportional not just to a global parameter but also to the amount of assets liquidated by the defaulted bank; we keep this refinement for future work.

## 4.2 Example of executions

First off, we remark that since all the actors in an ARS run in parallel, all the banks that have not been able to absorb the initial shock execute in parallel their default procedures. This might in principle lead to interferences between the behaviors of actors. As an example, if the bank  $b_2$  in Fig. 7 is not able to absorb the initial shock, it might receive the contagion message from Loan  $L_2$  while it is still dealing with its initial default. However, no interferences actually happen, because actors are implemented so that an incoming message is immediately enqueued in the actor's mailbox but the actor reacts to that message only after completing the reaction to the previously received message. In other terms, in ARSs the behavior of agents is resilient to a concurrent scenario where many systemic events happen in parallel. Then, assuming that  $b_2$  is not able to absorb the initial shock, in a system's execution the bank  $b_2$  receives the notification of the initial shock either before or after the contagion message from  $L_2$  (the initial default of  $b$  might be handled much more quickly). Accordingly, an inspection of the system's run will display the causal chains of events-and-reactions, thus witnessing either the default of  $b_2$  because of the initial shock or because of the

interbank contagion. Furthermore, we explicitly affected the message interleavings by letting the defaulting bank delay (by sleeping for a while) the notification to the Controller (notice that in Fig. 7 the defaulting bank  $b$  sleeps for a while). As a consequence, we let the contagion spread more quickly through the interbank contagion channel than through the global liquidity channel. The length of this delay is a model parameter so that it can be tuned to test different scenarios.

To exemplify, consider a run of the described European ARS with a severe initial shock and limited market illiquidity, i.e.,  $A = 0.2$  and  $P = 0.99$ . In this scenario the global result is that about 80 banks default, but the behavior of every bank can be precisely analyzed. The following is the log of a medium-sized Austrian bank that has been extracted from the complete system's log:

```

....
ATmed6 absorbed a shock of 395.0 caused by [init]
ATmed6 absorbed a shock of 78.9 caused by [ps-FRbig5,[init]]
ATmed6 absorbed a shock of 63.1 caused by [ps-ESbig1,[init]]
....
ATmed6 absorbed a shock of 17.5 caused by [cred-def ATmed3,[ps-FRbig5,[init]]]
ATmed6 absorbed a shock of 17.5 caused by [cred-def ATsm4,[ps-BEbig5,[init]]]
....
ATmed6 defaulted because [cred-def ATmed8,[cred-def ATmed13,[cred-def ATsm18,
                                                                    [ps-BEbig3,[init]]]]]
.....
ATmed4 absorbed a shock of 0.7 caused by [debt-def ATmed6,[cred-def ATmed8...]]
ATbig4 absorbed a shock of 157.9 caused by [cred-def ATmed6,[cred-def ATmed8..]
....
Global shock due to defaulted bank ATmed6
.... absorbed a shock of ... caused by [ps-ATmed6,[cred-def ATmed8 ...]]
IEmed4 defaulted because [ps-ATmed6,[cred-def ATmed8 ...]]
.....

```

The log shows that the Austrian bank absorbed the initial asset shock, as well as the global shocks due to the fire sales of the assets of two big foreign banks,  $FRbig5$  and  $ESbig1$ , that did not survive the initial shock. Moreover,  $ATmed6$  absorbed also the interbank contagions due to the default of its creditors  $ATmed3$  and  $ATsm4$ , which in turn did not survive to the fire sales caused by the two large foreign banks. We omitted the other shocks absorbed by the bank  $ATmed6$ , observing that even if these shock are absorbed, the bank's capital buffer is eroded, and indeed also  $ATmed6$  finally defaults. The cause of its default is a chain of defaulted creditors within the Austrian banking system: the small bank  $ATsm18$  did not survive to the global price shock caused by the initial default of  $BEbig3$ , then  $ATmed13$  did not survive to the default of its creditor  $ATsm18$ , then also  $ATmed8$  did not survive to the default of its creditor  $ATmed13$ , finally causing the default of  $ATmed6$  which was a debtor of  $ATmed8$ . We remark that, because of the nondeterministic interleaving of the messages that notify a new global fire sale, the name of the banks causing the first global fire sales may vary a lot in different system runs. However, the point is that  $ATmed6$  is able to absorb a number of global price shocks (of size  $P = 0.99$ ) but finally defaults because of an interbank contagion within its country. Also observe that the default of  $ATmed6$  causes another round of contagion and distress propagation. The log extract shows for instance that  $ATmed4$ , resp.  $ATbig4$ , absorbs the contagion due to the default of its debtor, resp. creditor, while the Irish bank  $IEmed4$  did not survive the global price shock caused by the default of the Austrian bank.

Further paradigmatic examples are taken from a run of the European ARS with moderate initial shock and severe market illiquidity, i.e.,  $A = 0.5$  and  $P = 0.2$ . The following long chains of internal contagions appear:

```
NLbig5 defaulted because [cred-def NLbig3, [cred-def NLmed9, [debt-def NLmed7,
                                                                    [init]]]]
DEmed5 defaulted because [cred-def DEmed10, [debt-def DEmed7, [cred-def DEbig1,
                                                                    [cred-def DEbig3, [cred-def ITbig3, [init]]]]]]
```

while the following is an example of external contagion chain exiting from The Netherlands and affecting Spain:

```
ESbig1 defaulted because [cred-def ESbig5, [cred-def NLbig5, [cred-def NLbig3,
                                                                    [cred-def NLmed9, [debt-def NLmed7, [init]]]]]]
```

To conclude, we observe that the fine-grained examination of the system's execution provided by the ARS is particularly suited to study the local properties of the system, and the contribution of individual components to the global dynamics. It allows, for instance, to identify those banks in the system that are more susceptible to interbank contagions compared to those that are more vulnerable to the market illiquidity. Furthermore, appropriate tests can reveal specific fragile interbank connections and typical sources of contagion within each country.

## 5 Analysis of the stress tests

In this section we report on the stress tests we conducted on the European ARS using the FinMarkSim tool. An exhaustive and in-depth analysis of the European banking network, providing financial insights on the systemic risk of the Eurosystem is postponed to future work. This will take advantage of the verification techniques available in the theory and practice of distributed systems (e.g., the model checking approach (Clarke et al. (2018); Baier and Katoen (2008)) is widely used in ICT), and by relying on a more experienced parameter calibration. Indeed, according to what happens in the related literature, the system behavior is significantly influenced by the parameters of the model.

As anticipated, by building an ARS directly as a software written in a highly scalable programming language, the system executions are so fast that it is very easy to develop stress tests in a test-oriented style, that is rapidly experimenting with many hypothetical scenarios and precisely tuning the model parameters. The behavior of the European ARS was tested under various conditions, so to investigate the role of the different parameters in the systemic risk. The risk is examined in terms of a number of systemic effects: the number and the size of defaulted banks, the depth and the main sources of the contagion, the volume of depositor losses. Each stress test has been run multiple times so to observe different executions due to the intrinsic nondeterminism of the actors' behavior. However, in all cases we observed that different runs of the same stress test showed almost identical global results (i.e., almost the same number of defaulted banks and almost the same depth of contagion) while differences appear in the details of specific contagion cascades.

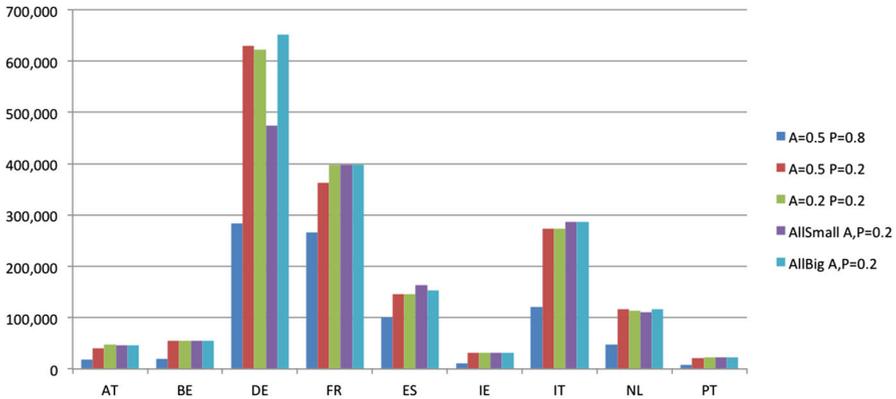


Fig. 8 Capital losses (in EUR Millions) with different sizes of initial shock and price shock

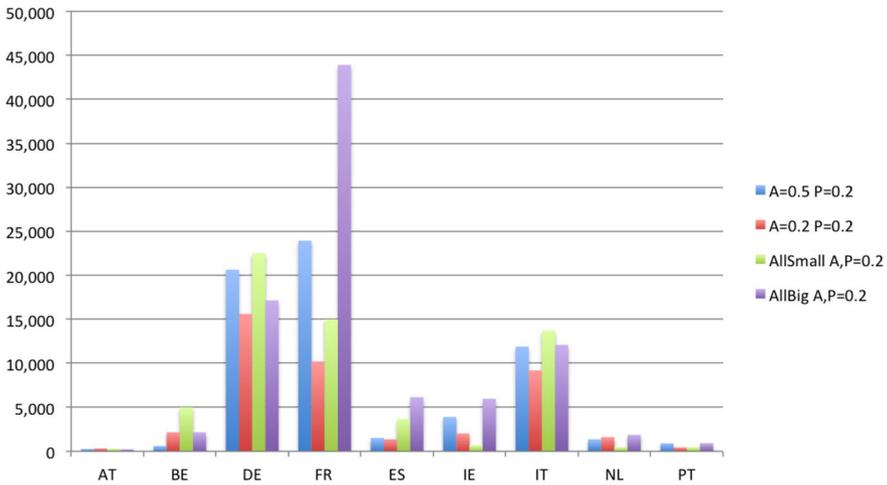
	A=0.5 P=0.8	A=0.5 P=0.2	A=0.2P=0.2	AllSmall A,P=0.2	AllBig A,P=0.2
AT	25,0	52,3	56,8	54,5	54,5
BE	27,3	68,2	68,2	68,2	68,2
DE	36,4	79,5	77,3	79,5	81,8
FR	31,8	45,5	50,0	50,0	50,0
ES	39,4	60,6	60,6	69,7	63,6
IE	22,7	59,1	59,1	59,1	59,1
IT	33,3	75,8	75,8	78,8	78,8
NL	27,3	66,7	66,7	63,6	66,7
PT	21,2	60,6	72,7	72,7	66,7
Capital loss	875.770	1.676.490	1.709.190	1.590.530	1.761.180
	36	69	70,3	65,5	72,5

Fig. 9 Percentages of defaulted bank in each country in different stress tests. The last two rows show the total capital loss in the European network (in EUR Millions) together with the corresponding percentage, given that the sum of the equities of all the European banks is 2.429.670 EUR Millions

### 5.1 Influence of the sizes of the initial shock and the market illiquidity

We observed that with low market illiquidity ( $P = 0.8$ ), a limited initial asset shock ( $A = 0.8$ ) is absorbed by all the banks. With a moderate initial shock ( $A = 0.5$ ) about 30% of banks fail. The global system shows to be rather robust in this condition of liquid market, since even with a radical initial shock ( $A = 0.005$ ) no more than 60% of the system collapses. With severe market illiquidity ( $P = 0.2$ ), again a limited initial shock ( $A = 0.8$ ) is absorbed by all the banks; however, a moderate initial shock ( $A = 0.5$ ) is enough to cause the default of about 64% of the system and the loss of 69% of the Eurozone capital. Nevertheless, even with drastic asset shock and market illiquidity ( $A, P = 0.005$ ), 80 banks of small and medium size still survive. This is in line with the fact that the contagion procedure is so that when assets are eroded, the role of the illiquidity contagion channel becomes limited.

In Figs. 8, 9 and 10 we show the percentage of defaulted banks, the capital lost and the depositor losses suffered in different countries in a number of stress tests.

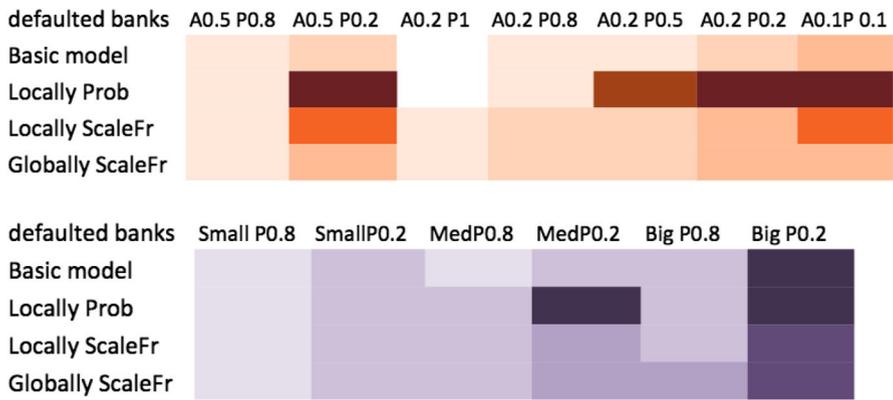


	A=0.5 P=0.2	A=0.2 P=0.2	AllSmall A,P=0.2	AllBig A,P=0.2
AT	6,8	4,5	4,5	4,5
BE	13,6	18,2	18,2	27,3
DE	15,9	20,5	20,5	27,3
FR	13,6	9,1	9,1	13,6
ES	15,2	15,2	18,2	21,2
IE	22,7	18,2	13,6	31,8
IT	24,2	30,3	30,3	30,3
NL	12,1	21,2	9,1	15,2
PT	24,2	30,3	30,3	36,4

**Fig. 10** Depositor losses (in EUR Millions) and percentages of banks that in each country suffered depositor losses in different stress tests

Figures 11, 12 emphasize the number of defaulted banks and the contagion depth caused by different stress tests and in different network models. As expected, by increasing the severity of the shocks, that is by decreasing the parameters  $A$  and  $P$ , both the number of defaulting banks and their contagion depth increase (up to contagion chains of length six and seven). However, when the market turmoil is beyond a certain level, the contagion depth does not further increase. Indeed, many banks immediately default because of the strong initial shock or because of few asset fire sales, thus leaving less room for contagion spreading. In general, we observed that in Austria, Germany and Portugal banks seems more susceptible to the internal contagion, exhibiting longer chains of interbank contagions between banks of the same country. Moreover, large banks are the sources of interbank contagions coming from foreign countries.

The stress test with parameters  $A = 0.2$  and  $P = 1$  is used to assess the role of a severe initial shock with no price shock due to the fire sale. In other terms, it allows one to observe just the interbank contagion by disallowing the illiquidity channel. In this case less than 90 banks default: large banks mostly default because of the severe asset shock, and a limited interbank contagion spreads to medium and small-sized local banks.



**Fig. 11** At the top: heatmap showing the number of *banks defaulted* in different stress tests and different models (the colors correspond to the intervals [0, 95), [95, 160), [160, 190), [190, 210), [210, 250), [250, 286]). At the bottom, heatmap showing the number of *banks defaulted because of a contagion* occurring by initially imposing the default of specific group of banks with different price-shock parameters (the colors correspond to the intervals: [0, 50), [50, 100), [100, 150), [150, 180), [180, 195), [195, 286])

### 5.2 Systemic influence of the immediate default of a country

In order to assess the systemic impact of a specific country in the Eurozone, we conducted a number of tests that impose the immediate default of specific groups of banks. Perhaps surprisingly, we observed that the default of a large country like Germany or France has a systemic influence that is not remarkably greater than that of countries with smaller banking sectors, like Portugal. In particular, we observed that most of the contagions progress externally via the liquidity channel: many failures are due to the price shock caused by the (initially imposed) default of some bank belonging to the country whose failure is under testing. The main difference is that, when defaulting a large country like Germany, a deeper interbank contagion spreads through the Eurozone. Instead, by imposing the initial default of a smaller country like Portugal, the illiquidity channel assumes a more relevant role in the systemic contagion.

We also observed that the systemic effect of initially defaulting the banks of some country with price shock parameter  $P = p$  is approximately the same that is obtained by shocking all the European system with parameters  $A, P = p$ . Therefore, we think that the model of the European interbank network considered in this paper, i.e., the balance sheets structure, the edge generation and the contagion procedure, emphasize the role of the market illiquidity over the that of the failures of creditors and debtors as the source of contagion. A deeper analysis of the assumptions used in this model is an important subject of future work; however, we observe that the ARSs’ ability to precisely disentangle and observe the roles of the interbank and illiquidity channels allowed to make interesting observation about the contagion spreading.

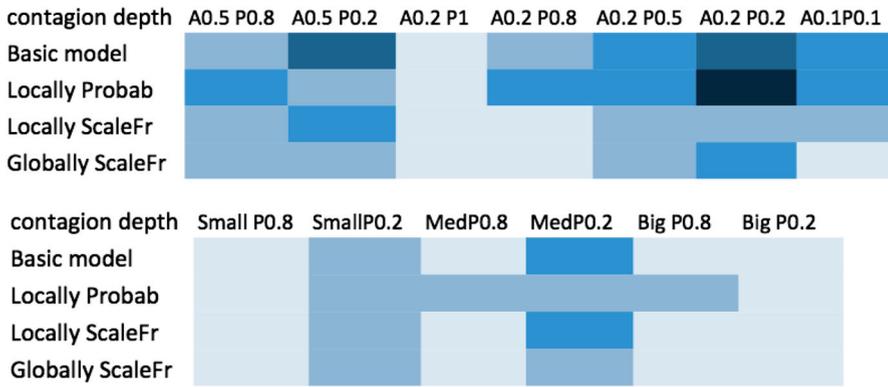


Fig. 12 Depth level of the contagion chains in different stress tests and in different models

### 5.3 Influence of the size of banks

As illustrated in Figs. 11 and 12, the size of banks shows a systemic role. The immediate default of all the small banks in the network causes fewer global defaults than the failure of larger banks. By imposing the default of medium-sized banks, we observe that small banks tend to absorb the default of medium-seized creditors and debtors, but fail because of the global price shock. On the other hands, large banks tend to fail because of (possibly long) chains of interbank contagions due to the default of large or medium-sized (foreign) banks.

### 5.4 Influence of the graph topology

We investigated the role of node interconnectedness by running the same stress tests on the four different graph topologies described in Sect. 3. We remind that, except for the globally scale-free network, in all the other graphs the large foreign banks are connected with the same uniform probability, while local country graphs are built according to the different topologies.

We observed that in the locally probabilistic network the systemic effect of every stress test is amplified (see Figs. 11 and 12). This is in line with the fact that this network contains many more edges and all the nodes are much more interconnected than in the other graphs (see Fig. 3). The locally scale-free graph behaves very similarly to the default network, providing evidence about the fact that our method of connecting banks of different sizes using different probabilities actually results in a graph that is roughly scale-free.

Finally, the important role of central nodes (also called “too-connected-tot-fail”) is confirmed by the behavior of the globally scale-free graph. Indeed, in this network the node connectedness is proportional to the bank size. Accordingly, Fig. 11 shows that the initial default of larger banks causes the default of a larger set of (other) banks.

## 6 Conclusions

We illustrated a computational model, called Actor-based Reactive System, that similarly to Agent-Based Models, studies the system's behavior in terms of the interactions of its local agents. However, ARSs depart from the established approach of ABMs in two ways: a fundamentally different management of time, and a fully decentralized control logic, that have a profound impact on the expressiveness of the simulations. Moreover, the highly scalable implementation allows to effectively deal with large financial systems with a higher degree of agents' heterogeneity, and to develop very flexible stress tests to rapidly experiment with many hypothetical scenarios in a test-oriented style.

We detailed the construction and the analysis of a realistic case study by considering an ARS that models the systemic risk of the European banking network with a nontrivial contagion procedure, that combines an initial asset shock with the negative feedback loop triggered by asset fire sales. We showed that the asynchronous and decentralized implementation of the agents' interactions brings to the fore the interleaving of events, exposing chains of causally related events. This allows to precisely disentangle and assess the roles of the two contagion channels operating in the system, that is the interbank contagion and the illiquidity channel.

We plan for future work the development of an in-depth analysis of the European ARS, in order to provide financial insights and policy recommendations about the systemic risk of the European interbanking system. We also plan to exploit the flexibility of the architectural pattern to enrich the model in various ways, for instance by letting the banks borrow extra funds during the global shock phase, thus dynamically adding new edges to the interbank network. This would allow to model appropriate recovery strategies while the system is in distress, such as letting a bank increment its capital position before being reached by the contagion. Other interesting extensions are the generalization to different kinds of assets and the inclusion of financial derivatives like Credit Default Swaps, as in the contagion procedure studied in Crafa and Varacca (2018). We remark that all these extensions can still be obtained as an instance of the general pattern; indeed, they do not require a global change of the ARS, but only the addition of new classes of actors and new sets of messages (with corresponding reactions) to be exchanged between the economic agents.

To conclude, we observe that, compared to ABM, the fine-grained inspection of the system's execution provided by the ARS framework is particularly suited to the study of financial systems, that are characterized by individual components that are quite heterogeneous. Indeed, compared to biological or physical systems, that are made of a multitude of homogeneous components, in financial systems each economic agent has a clear identity, that sets him apart from the others. Moreover, every financial institution holds private information that is not disclosed but that nevertheless influence its behavior. In this kind of systems, a specific agent might intentionally deviate from the average behavior, thus it is of particular interest to be able to observe and assess its specific role within the global system's dynamics.

**Funding** Open access funding provided by Università degli Studi di Padova within the CRUI-CARE Agreement.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Agha G (1986) *Actors: a model of concurrent computation in distributed systems*. MIT Press, Cambridge
- Anand K, van Lelyveld I, Banai A, Friedrich S, Garratt R, Halaj G, Figue J, Hansen I, Jaramillo SM, Lee H, Molina-Borboa JL, Nobili S, Rajan S, Salakhova D, Silva TC, Silvestri L, de Souza SRS (2018) The missing links: a global study on uncovering financial network structures from partial data. *J Financ Stab* 35:107–119
- Baier C, Katoen J (2008) *Principles of model checking*. MIT Press, Cambridge
- Battiston S, Puliga M, Kaushik R, Tasca P (2012) Debrank: too central to fail? financial networks, the fed and systemic risk. *Sci Rep* 2(541):1–6
- Battiston S, Caldarelli G, D'Errico M, Gurciullo S (2016) Leveraging the network: a stress-test framework based on debrank. *Stat Risk Model* 33(3–4):117–138
- BCBS (2015) Making supervisory stress tests more macroprudential: Considering liquidity and solvency interactions and systemic risk. Basel Committee on Banking Supervision's Research Task Force Working Paper 29
- Capponi A (2016) Systemic risk, policies, and data needs. *INFORMS Tutorials Oper Res*, pp 185–206. <https://doi.org/10.1287/educ.2016.0156>
- Capponi A, Bo L (2015) Systemic risk in interbanking networks. *SIAM J Financ Math* 6(1):386–424
- Cimini G, Squartini T, Saracco F, Garlaschelli D, Gabrielli A, Caldarelli G (2019) The statistical physics of real-world networks. *Nat Rev Phys* 1:58–71
- Clarke EM, Grumberg O, Kroening D, Peled D, Veith H (2018) *Model checking*, 2nd edn. MIT Press, Cambridge
- Cont R, Minca A (2016) Credit default swaps and systemic risk. *Annals Op Res* 247:523–547
- Cont R, Moussa A, Santos E (2013) *Network structure and systemic risk in banking systems*. Cambridge University Press, Cambridge, pp 327–368
- Crafa S, Varacca D (2018) On the reactive nature of financial networks. In: 19th Italian Conference on Theoretical Computer Science, 2018, CEUR-WS.org, vol 2243, pp 230–241, <http://ceur-ws.org/Vol-2243/paper23.pdf>
- Dai Pra P, Runggaldier W, Sartori E, Tollot M (2009) Large portfolio losses: a dynamic contagion model. *Annals Appl Prob* 19(1):347–394
- De Masi G, Iori G, Caldarelli G (2006) Fitness model for the Italian interbank money market. *Phys Rev E* 74(6):066112
- ECB (2018) Aggregated balance sheet of euro area credit institutions. <http://sdw.ecb.europa.eu/reports.do?node=1000005719>
- Farmer J, Foley D (2009) The economy needs agent-based modelling. *Nature* 460:685–686
- Gandy A, Veraart LAM (2019) Adjustable network reconstruction with applications to cds exposures. *J Multivar Anal* 172:193–209
- Iori G, De Masi G, Precup O, Gabbi G, Caldarelli G (2008) A network analysis of the Italian overnight money market. *J Econ Dyn Control* 32(1):259–278
- Neveu A (2018) A survey of network-based analysis and systemic risk measurement. *J Econ Interact Coord* 13:241–281
- Ramadhani A, Caccioli F, Fricke D (2020) Reconstructing and stress testing credit networks. *J Econ Dyn Control* 111:103817
- Seppelcher P (2012) Flexibility of wages and macroeconomic instability in an agent-based computational model with endogenous money. *Macroecon Dyn* 16:284–297

- Soramaki K, Bech M, Arnold J, Glass R, Beyeler W (2007) The topology of interbank payment flows. *Physica* 397:313–333
- Teply P, Klinger T (2019) Agent-based modeling of systemic risk in the european banking sector. *J Econ Interact Coord* 14:811–833
- Tesfatsion L, Judd K (eds) (2006) *Handbook of computational economics. agent-based computational economics*. North Holland, Elsevier
- Trichet JC (2010) Opening address : reflections on the nature of monetary policy non-standard measures and finance theory. In: Bank EC (ed) 6th ECB Central Banking Conference, vol 18

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.