

Financial Networks with Singleton Liability Priorities

Stavros D. Ioannidis, Bart de Keijzer, and Carmine Ventre^(\boxtimes)

King's College London, London, UK {stavros.ioannidis,bart.de_keijzer,carmine.ventre}@kcl.ac.uk

Abstract. Financial networks model debt obligations between economic firms. Computational and game-theoretic analyses of these networks have been recent focus of the literature. The main computational challenge in this context is the clearing problem, a fixed point search problem that essentially determines insolvent firms and their exposure to systemic risk, technically known as recovery rates. When Credit Default Swaps, a derivative connected to the 2008 financial crisis, are part of the obligations and insolvent firms pay the same proportion of all their debts, computing a weakly approximate solution is PPAD-complete [29], whereas computing a strongly approximate solution is FIXP-complete [17]. This paper addresses the computational complexity of the clearing problem in financial networks with derivatives, whenever priorities amongst creditors are adopted. This practically relevant model has been only studied from a game-theoretic standpoint. We explicitly study the clearing problem whenever the firms pay according to a singleton liability priority list and prove that it is FIXP-complete. Finally, we provide a host of NP-hardness results for the computation of priority lists that optimise specific objectives of importance in the domain.

1 Introduction

The financial services industry has been very creative, with the constant introduction of new products designed as investment and/or risk management tools. This makes the web of liabilities between the different institutions in the market hard to track and oversee. It is, in fact, this inherent complex structure of the evolving modern financial system that has led to several somewhat unforeseen and deeply damaging crises, such as the Great Financial Crisis (GFC) of 2008. There is, therefore, the need to mathematically model and study this network of obligations among interconnected financial agents in order to understand the impact of new products, regulations or even single contracts.

The main computational challenge in this context is the *clearing problem* introduced in [6]: Given the banks' funds and the face values of all the liabilities in the network, compute for each bank its exposure to systemic risk, in the form of what is known as its *clearing recovery rate*. It turns out that the

C. Ventre—Partially supported by the UKRI grant EP/V00784X/1.

[©] The Author(s), under exclusive license to Springer Nature Switzerland AG 2022 P. Kanellopoulos et al. (Eds.): SAGT 2022, LNCS 13584, pp. 205–222, 2022. https://doi.org/10.1007/978-3-031-15714-1_12

complexity of this problem is closely related to the class of financial products populating the network. In fact, in financial networks where firms only subscribe simple debt contracts, clearing recovery rates can be computed in polynomial time [6,25]. Whilst this setup makes up for an easy modeling of networks as directed graphs with nodes standing for financial institutions and edges representing debt obligations, this representation is too simplistic in that it does not capture more advanced products, based on other existing contracts like mortgages, loans, interest rates etc. These complex contracts are called *derivatives*. The addition of one such derivative, namely, Credit Default Swaps (CDSes), introduced to financial networks in [28], makes the model more intriguing from a computational perspective. A CDS involves three parties i, j, k, where k must pay an amount of money to i on the condition that i cannot pay off all of its obligations. Since adding CDSes to a financial network may generate irrational clearing recovery rates [17,28], our interest turns to finding approximate clearing recovery rates; it is proved that weakly (or "almost") approximate recovery rates are PPAD-complete to compute [29] and strong (or "near") approximate solutions are FIXP-complete to compute [17].

In this paper, we look at this problem from the perspective of the financial regulator. We ask whether rules that determine how insolvent banks pay off their debts can fundamentally change the computational hardness landscape above. The most-studied payment scheme is the *proportional payment scheme* where each bank pays off its debts proportionally to its recovery rate. However, defining priority classes amongst creditors and pay proportionally in each class (with funds available at the current priority level) is another widely adopted measure used in practice in the industry. For example, some regulatory regimes require employees to be prioritised over other creditors whereas some advanced derivatives (such as, the renown Collateralized Debt Obligations leading to the GFC) define their payoff via *tranching* (effectively a priority list) of the underlying securities. Whilst priority payments have been studied under a game-theoretic framework [3,18,21], nothing is known about the complexity of the clearing problem with this payment scheme in presence of financial derivatives.

Our Contribution. We study the clearing problem in financial networks with derivatives under the *priority list* payment scheme. Specifically, we examine financial networks consisting of both debt and CDS contracts and address the complexity of computing a *clearing recovery rate vector* whenever each bank has to pay its debts in the following way. For each bank, we define a partition of its liabilities in priority classes. With the funds available at a certain priority (i.e., after having paid all the liabilities with higher priority), the liabilities at the current priority are paid proportionally – in particular, this means that these are paid in full if the funds are sufficient. This notion generalises the proportional payment scheme studied in related literature (i.e., consider the case in which the partition contains one part) and the class of *singleton priorities*, where each part is a singleton. We call this problem CDS-PRIORITY-CLEARING. Note that without CDSes, CDS-PRIORITY-CLEARING is known to be in P, both for proportional [6] and priority [18] payment schemes.

We observe here that whenever the partition defined by the priority list contains at least of part of size 2 or more then CDS-PRIORITY-CLEARING is FIXPcomplete. FIXP [8] is a complexity class that captures the fixed point computations of total search problems. In our context, it is important to observe that FIXP is defined in terms of fixed-point functions defined upon the algebraic basis $A = \{+, -, *, /, \max, \min, \frac{k}{2}\}$. It is not hard to see that the FIXP-completeness follows from the recent reduction given in [17] where all gadgets adopted have maximum out-degree 2.¹ Therefore, the only case left open is when all the parts of the priority list are singletons. We focus our attention on this setup in our work. We then call CDS-PRIORITY-CLEARING the problem of finding recovery rates for singleton priorities, i.e., each bank has an ordering of its liabilities according to which its debts are paid off.

Our first contribution gives technical evidence that the financial regulator cannot change the complexity of computing clearing recovery rates, by enforcing priority payments. Specifically, we prove that CDS-PRIORITY-CLEARING is at least as hard as the square root sum (SQRT-SUM) problem [8,11,20,30] and that it is complete for FIXP. We then give the full picture of the complexity for the clearing problem, under all payment schemes proposed in the literature. Whilst the proof of FIXP-completeness adopts known approaches [17], our reduction introduces new financial network gadgets with priority payments for the operations in A. In particular, our new multiplication gadget highlights the flexibility of financial networks in handling arithmetic operations.

Whilst the regulator cannot ease the computational complexity of the problem, we wonder whether one can efficiently compute the banks' priority lists to optimise certain objective functions of financial interest. These include maximising the equity of a specific bank, maximising the liquidity in the system, and minimising the number of activated CDSes. As our second main contribution, we present a set of NP-hardness results showing a parallel with the known hardness of computing similarly "optimal" solutions with proportional payments [22].

Related Work. Clearing problems have been studied a lot in the literature [1,5–7,12,14–16,25]. Analysis of financial networks with CDSes as well as their properties is a popular topic in the area [17,21,23,24,27–29]. A game-theoretic approach to financial networks, the edge ranking game as well as other financial network games are listed in [3,18,19,22]. The FIXP-complexity class was defined first in [8,17,31], which established the FIXP-completeness of various fundamental fixed point computation problems. There are various further recent papers that show FIXP-completeness of a range of problems, including [9,10,13,17].

2 Model and Preliminaries

Financial Networks and Payment Schemes. A financial network consists of a set of financial entities (which we refer to as *banks* for convenience), interconnected through a set of financial contracts. Let $N = \{1, ..., n\}$ be the set

¹ We defer a formal treatment of this and the omitted proofs herein to the full version.

of n banks. Each bank $i \in N$ has external assets, denoted by $e_i \in \mathbb{Q}_{\geq 0}$. We let $e = (e_1, \ldots, e_n)$ be the vector of all external assets. We consider two types of liabilities among banks: debt contracts and credit default swaps (CDSes). A debt contract requires one bank i (debtor) to pay another bank j (creditor) a certain amount $c_{i,j} \in \mathbb{Q}_{\geq 0}$. A CDS requires a debtor i to pay a creditor j on condition that a third bank called the reference bank R is in default, meaning that R cannot fully pay its liabilities. Formally, we associate each bank i with a recovery rate $r_i \in [0, 1]$, that indicates the proportion of liabilities i is able to pay. Having $r_i = 1$ means that bank i can fully pay its liabilities, while $r_i < 1$ indicates that i is in default. In case a reference bank R of a CDS is in default, the debtor i of that CDS pays the creditor j an amount of $(1 - r_R)c_{i,j}^R$, where $c_{i,j}^R \in \mathbb{Q}_{\geq 0}$ is the face value of the CDS. The value $c_{i,j}$ ($c_{i,j}^R$, resp.) of a debt contract (CDS, resp.) is also called the notional of the contract. Finally, we let c be the collection of all contracts' notionals. We do not allow any bank to have a debt contract with itself, and assume that all three banks in any CDS are distinct.

The financial system \mathcal{F} can therefore be represented as the triplet (N, e, c). Given \mathcal{F} , we let $\mathcal{DC}_{\mathcal{F}}$ denote the set of all pairs of banks participating in a debt contract in \mathcal{F} . Similarly, $\mathcal{CDS}_{\mathcal{F}}$ denotes the set of all triplets participating in a CDS in \mathcal{F} . (We drop \mathcal{F} from the notation when it is clear from the context.) The contract graph of $\mathcal{F} = (N, e, c)$ is defined as a coloured directed multigraph $G_{\mathcal{F}} = (V, A)$, where V = N and $A = (\bigcup_{k \in N} A_k) \cup A_0$ where $A_0 = \{(i, j) \mid i, j \in N \land c_{i,j} \neq 0\}$ and $A_k = \{(i, j) \mid i, j \in N \land c_{i,j}^k \neq 0\}$. Each arc $(i, j) \in A_0$ is coloured blue and each $(i, j) \in A_k$ orange. For all $(i, j, R) \in \mathcal{CDS}$ we draw a dotted orange line from node R to arc $(i, j) \in A_R$, denoting that R is the reference bank of the CDS between i and j.² Finally, we label each arc with the notional of the corresponding contract, and each node with the external assets of the corresponding bank.

All banks are obliged to pay off their liabilities according to a set of prespecified rules, which we refer to as a payment scheme. If a bank has sufficient assets, then the payment scheme is trivial and prescribes to simply make payments that correspond exactly to each of the bank's liabilities. If there are insufficient assets, on the other hand, the payment scheme will determine for each of the outgoing contracts how much of it is paid off. The most studied payment scheme is the proportional payment scheme, where each bank i submits an r_i proportion of each liability, leaving a $(1-r_i)$ fraction of each liability unpaid. In this paper we study payments resulting from another rule, called the singleton liability priority lists payment scheme. More specifically, given a financial system $\mathcal{F} = (N, e, c)$, for each bank i we define a total order over the arcs going out of i in $G_{\mathcal{F}}$. We denote the singleton liability priority list of i as $P_i = (i_1 \mid i_2 \mid ... \mid i_{outdeg(i)})$, where i_k stands for the k-th element in the order, or kth priority of node i, and outdeg(i)denotes the out degree of node i in $G_{\mathcal{F}}$. The payments under this scheme are now formed through an iterative process where each bank pays off its liabilities, one after the other, preserving the ordering in its priority list. We denote by c_{i_k} the contract notional of the i_k th priority and denote by $\mathcal{P} = (P_1, \ldots, P_n)$ the profile

² This means that $G_{\mathcal{F}}$ is a directed hypergraph with arcs of size 2 and 3.

of singleton liability priority lists. We denote a financial system \mathcal{F} endowed with a singleton liability priority profile \mathcal{P} as $(\mathcal{F}, \mathcal{P})$. The next example illustrates the model.

Example 1. The financial system of Fig. 1 consists of six banks, $N = \{1, 2, 3, 4, 5, 6\}$. Banks 2 and 5 have external assets $e_2 = e_5 = 1 - c$, for some constant $c \in (0, 1)$, while all other banks have zero external assets. The set of debt contracts is $\mathcal{DC} = \{(2, 3), (5, 4)\}$ and the set of CDS contracts is $\mathcal{CDS} = \{(2, 1, 5), (5, 6, 2)\}$. All contract notionals are set to 1. For example, $c_{2,3} = c_{2,1}^5 = 1$. Node 2 has two candidate singleton liability priority lists, one is $P_2^1 = ((2, 3) \mid (2, 1, 5))$, where $2_1 = (2, 3)$ with contract notional $c_{2_1} = c_{2,1} = 1$ and $2_2 = (2, 1, 5)$ with contract notional $c_{2_2} = c_{2,1}^5 = 1$. The other one is $P_2^2 = ((2, 1, 5) \mid (2, 3))$ where $2_1 = (2, 1, 5)$ with $c_{2_1} = c_{2,1}^5 = 1$ and $2_2 = (2, 3)$ with $c_{2_2} = c_{2,3} = 1$. Symmetrically one can derive the lists for node 5.



Fig. 1. Example of a financial network

We are interested in computing for a pair $(\mathcal{F}, \mathcal{P})$, for each bank i, the proportion of liabilities that it is able to pay. This proportion is captured by the *recovery rate*, mentioned earlier: For each bank i we associate a variable $r_i \in [0, 1]$, that indicates the proportion of liabilities that bank i can pay. Recall that, to define the liability generated from a CDS contract, we need the recovery rate of the reference banks. Consequently in order to define all liabilities of banks in a financial system, we need to be presented with an a-priori recovery rate vector $r = (r_1, \dots, r_n)$. So given a $(\mathcal{F}, \mathcal{P})$ and assuming a vector $r \in [0, 1]^n$, we define the liabilities, the payments that each bank submits and the assets for each bank as follows. We denote by $l_{i_k}(r)$ the k-th liability priority of node i. If $i_k = (i, j) \in \mathcal{DC}$ for some $j \in N$, then $l_{i_k}(r) = c_{i,j}$ and if $i_k = (i, j, R) \in \mathcal{CDS}$ for some $j, R \in N$, then $l_{i_j}(r) = (1 - r_R)c_{i,j}^R$. The liability of bank $i \in N$ to a bank $j \in N$ is denoted by $l_{i,j}(r)$ and it holds that

$$l_{i,j}(r) = c_{i,j} + \sum_{k \in N} (1 - r_k) c_{i,j}^k.$$
 (1)

We denote by $l_i(r)$ the total liabilities of node *i*, and it holds that

$$l_i(r) = \sum_{j=1}^{\text{outdeg}(i)} l_{i_j}(r) = \sum_{j \neq i} l_{i,j}(r) = \sum_{j \in N \setminus \{i\}} \left(c_{i,j} + \sum_{k \in N \setminus \{i,j\}} (1 - r_k) c_{i,j}^k \right).$$

Node *i* can fully pay its *k*th priority only if he has sufficient assets left after paying off the liabilities corresponding to priorities i_1, \ldots, i_{k-1} . We denote by $p_{i_k}(r)$ the payment of node *i* to its *k*-th priority, and by $a_i(r)$ its assets, which are defined as the external assets it possesses plus all incoming payments submitted from its debtors (see below for a more formal definition). Under our singleton priority lists payment scheme:

$$p_{i_j}(r) = \max\left\{0, \min\left\{l_{i_j}(r), a_i(r) - \sum_{j' < j} l_{i_{j'}}(r)\right\}\right\}.$$
(2)

Moreover, we denote by $p_{i,j}(r)$ the payment of node *i* to node *j* under recovery rate vector *r*: Let $C_j = \{i_j \mid i_j \text{ is a contract with node$ *j* $as the creditor}\}$, then $p_{i,j}(r) = \sum_{i_j \in C_j} p_{i_j}(r)$. The total payment made by a node is the sum of its individual payments to its priorities which is equal to the total sum of its payments to its creditors. Also, the payment of node *i* needs to be equal to the proportion of its total liabilities it can pay off. Therefore, the following equations hold.

$$p_i(r) = \sum_{j=1}^{\text{outdeg}(i)} p_{i_j}(r) = \sum_{j \in N} p_{i,j}(r) = r_i l_i(r).$$
(3)

The assets of a bank i are denoted as $a_i(r)$ and are the total amount of money it possesses summing its external assets and all incoming payments made by its debtors. It holds that

$$a_i(r) = e_i + \sum_{j \neq i} p_{j,i}(r).$$
 (4)

We are interested in computing a specific recovery rate vector in $[0, 1]^n$, such that (3) holds (i.e., $p_i(r) = r_i l_i(r)$ for all $i \in N$), under the singleton liability priority list payment rule (just defined formally by (2) and (4)). Formally:

Definition 1 (Clearing recovery rate vectors (CRRVs)). Given a financial system and a singleton liability priority profile $(\mathcal{F}, \mathcal{P})$, a recovery rate vector r is called clearing if and only if for all banks $i \in N$,

$$r_i = \min\left\{1, \frac{a_i(r)}{l_i(r)}\right\} \text{ if } l_i(r) > 0, \text{ and } r_i = 1 \text{ if } l_i(r) = 0.$$
(5)

We illustrate the notions of the dynamics and the CRRVs by reconsidering Example 1 and computing them for some priority profile \mathcal{P} .

Example 1 (continued). Let c = 1/4 in Fig. 1. Let $\mathcal{P} = (P_2 = ((2,3) \mid (2,1,5)), P_5 = ((5,4) \mid (5,6,2)))$. Both nodes 2 and 5 receive no payment from any other node thus their assets are defined as $a_2 = e_2 = 1-c$ and $a_5 = e_5 = 1-c$. For node 2, given P_2 , we get that $l_{2_1} = l_{2,1} = c_{2_1} = c_{2,3} = 1$ and $l_{2_2} = l_{2,1} = (1-r_5)c_{2,1}^5 = (1-r_5)$, thus the total liabilities for node 2 are $l_2 = l_{2_1} + l_{2_2} = 2 - r_5$. For node 5 we get that $l_{5_1} = l_{5,4} = c_{5_1} = c_{5,4} = 1$ and $l_{5_2} = l_{5,6} = c_{5_2} = (1-r_2)c_{5,6}^2 = 1-r_2$, thus the total liabilities for node 5 are $l_5 = l_{5_1} + l_{5_2} = 2 - r_2$. Let us compute the CRRV. By (5) it must be $r_2 = \min\{1, a_2(r)/l_2(r)\} = \min\{1, 1 - c/2 - r_5\}$ and

 $r_5 = \min \{1, a_5(r)/l_5(r)\} = \min \{1, (1-c)/(2-r_2)\}$. After solving this system we get that $r_2 = r_5 = 1 - \sqrt{c}$ and since we assumed c = 1/4 we finally get that $r_2 = r_5 = 1/2$. For the payments of node 2, we know that $a_2 = 3/4$ and it first prioritises node 3 for which it has a liability of 1, thus it cannot fully pay off that liability and submits all of its assets to node 3, namely $p_{2_1} = p_{2,3} = 3/4$ and $p_{2_2} = p_{2,1} = 0$. The payments of node 5 are symmetrical.

Our Search Problem. We define CDS-PRIORITY-CLEARING to be the search problem that asks for a clearing recovery rate vector r given a pair $(\mathcal{F}, \mathcal{P})$. The term CDS refers to the fact that \mathcal{F} may contain CDS contracts (the problem becomes polynomial time computable without CDSes [18]) and the term PRI-ORITY indicates that banks pay according to singleton liability priority list \mathcal{P} . Similarly to [17,29], we assume that \mathcal{F} is non-degenerate (see below for a discussion).

Definition 2. A financial system is non-degenerate if and only if the following two conditions hold. Every debtor in a CDS either has positive external assets or is the debtor in at least one debt contract with a positive notional. Every bank that acts as a reference bank in some CDS is the debtor of at least one debt contract with a positive notional.

Given an instance $I \in \text{CDS-PRIORITY-CLEARING}$ we transform (1) into a function defined on arbitrary recovery rate vectors $r = (r_1, \dots, r_n) \text{ as:}^3$

$$f_I(r)_i = \frac{a_i(r)}{\max\{a_i(r), l_i(r)\}}.$$
(6)

From (6) we ascertain that r is a clearing recovery rate vector for I if and only if r is a fixed point of f_I , namely $r = f_I(r)$. Thus, solving CDS-PRIORITY-CLEARING comes down to computing the fixed points of f_I . We define CDS-PRIORITY-CLEARING to contain only non-degenerate financial networks, for the analytical convenience that non-degeneracy provides (note that a division by 0 never occurs in $f_I(r)_i$ for these instances). It is not hard to see that f_I has fixed points, see, e.g., [21]. Moreover, there exist instances of $(\mathcal{F}, \mathcal{P})$ that admit multiple CRRVs.

Irrationality. As is the case for the proportional payments, the singleton liability priority list model contains instances that admit irrational CRRVs. Observation 1 below provides such an example while Observations 2 and 3 present examples of how the priority profile affects the payments in the network. These examples also provide insights on an important difference between the two payment scheme: In the proportional model, whenever the CRRV is irrational then the clearing payment vector must be irrational as well. That is not the case in the singleton liability priority payment scheme, where irrationality of a CRRV need not cause any irrationality in the payments.

³ Strictly speaking, $f_I(r)_i$ is defined only for nodes *i* that are not sinks in the contract graph. Sink nodes have recovery rate 1, cf. (5). Their exclusion simply allows to bypass potential divisions by 0 in f_I (e.g., take node 1 in Fig. 1 when c = 1) while preserving its continuity. For notational simplicity, we will implicitly assume that we compute $f_I(r)_i$ iff *i* is not a sink.

Observation 1. There exist instances of $(\mathcal{F}, \mathcal{P})$ that have irrational CRRVs. For instance, we know that $r_2 = r_5 = 1 - \sqrt{c}$ in Example 1. Thus, it is clear that for many choices of $c \in (0, 1)$ (e.g., c = 1/2) the CRRV is irrational.

Observation 2. There exist a pair $(\mathcal{F}, \mathcal{P})$ with an irrational CRRV and irrational payments. Take again Example 1 and fix c = 1/3. We have $e_2 = e_5 = 2/3$, $l_{2,3} = l_{5,6} = 2/3$ and $r_2 = r_5 = 1 - \sqrt{1/3}$. Now consider the singleton liability priority lists $P_2 = ((2,1,5) \mid (2,3))$ and $P_5 = ((5,6,2) \mid (5,4))$. Since node 2 prioritises the (2,1,5) contract, it has to pay an amount of $1 - \sqrt{1/3}$ to node 1. Given that its total assets are 2/3, it can fully pay this liability and so $p_{2,1} = 1 - \sqrt{1/3}$ and what is left is being paid to node 3. Symmetrically, one can compute that $p_{5,6} = 1 - \sqrt{1/3}$.

Observation 3. There exist a pair $(\mathcal{F}, \mathcal{P})$ with an irrational CRRV and rational payments. Consider Example 1 once more and fix c = 1/2. This yields $e_2 = e_5 = 1/2$, $l_{2,3} = l_{5,6} = 1/2$ and $r_2 = r_5 = 1 - \sqrt{1/2}$. Consider the singleton liability priority lists $P_2 = ((2,1,5) \mid (2,3))$ and $P_5 = ((5,6,2) \mid (5,4))$. Since node 2 prioritises the (2,1,5) contract, it has to pay a amount of $1 - \sqrt{1/2}$ to node 1 but only possesses total assets of 1/2. Thus it cannot fully pay this liability, meaning that $p_{2,1} = 1/2$. Symmetrically, we can compute that $p_{5,6} = 1/2$.

A Primer on FIXP. A useful framework for studying the complexity of fixed point computation problems is defined in [8]. Both exact and approximate computation of the solutions to such problems are considered. We begin by defining the notion of approximation we are interested in. Let F be a continuous function that maps a compact convex set to itself and let $\epsilon > 0$ be a small constant. An ϵ approximate fixed point of F is a point x is within a distance ϵ near a fixed point of F, i.e., $\exists x' : F(x') = x' \land \|x' - x\|_{\infty} < \epsilon$. This notion is also known as strong approximation.⁴ We now introduce the problems we are focusing on. A fixed point problem Π is defined as a search problem such that for every instance $I \in \Pi$ there is an associated continuous function $F_I : D_I \to D_I$ —where $D_I \subseteq \mathbb{R}^n$ (for some $n \in \mathbb{N}$) is a convex polytope described by a set of linear inequalities with rational coefficients that can be computed from I in polynomial time—such that the solutions of I are the fixed points of F_I .

Definition 3. The class FIXP consists of all fixed point problems Π for which for all $I \in \Pi$ the function $F_I : D_I \to D_I$ can be represented by an algebraic circuit C_I over the basis $\{+, -, *, /, \max, \min, \sqrt[k]{}\}$, using rational constants, such that C_I computes F_I , and C_I can be constructed from I in time polynomial in |I|.

The class $FIXP_a$ is defined as the class of search problems that are the strong approximation version of some fixed point problem that belongs to FIXP.

⁴ A weak ϵ -approximate fixed point of a continuous function F is a point x such that its image is within distance ϵ of x, i.e., $||x - F(x)||_{\infty} < \epsilon$. Under polynomial continuity, a mild condition on the fixed point problem under consideration, a strong approximation is also weak [8].

The class Linear-FIXP is defined analogously to FIXP, but under the smaller arithmetic basis where only the gates $\{+, -, \max, \min\}$ and multiplication by rational constants are used.

These classes admit complete problems. The completeness results in [8] in fact show that it is without loss of generality to consider a restricted basis {+, *, max} ({+, max}) for FIXP (Linear-FIXP), and to assume that $D_I = [0, 1]^{n.5}$ Hardness of a search problem Π for FIXP is defined through the existence of a polynomial time computable function $\rho : \Pi' \to \Pi$, for all $\Pi \in \text{FIXP}$, such that the solutions of I can be obtained from the solutions of $\rho(I)$ by applying a (polynomial-time computable) linear transformation on a subset of $\rho(I)$'s coordinates. This type of reduction is known as a *polynomial time SL-reduction*.

It is known that $FIXP_a \subseteq PSPACE$ and Linear-FIXP = PPAD [8]. An informal understanding of the hardness of FIXP vis-a-vis PPAD is as follows. PPAD captures a type of computational hardness stemming from an essentially combinatorial source. FIXP introduces on top of that a type of numerical hardness that emerges from the introduction of multiplication and division operations: These operations give rise to irrationality in the exact solutions to these problems, and may independently also require the computation of rational numbers of very high precision or very high magnitude.

3 Hardness of CDS-PRIORITY-CLEARING

We are interested in identifying the complexity of CDS-PRIORITY-CLEARING. Recall that in Example 1, we presented an instance which under proper coefficients admits only irrational clearing recovery rate vectors, which means that either one should study this problem with respect to complexity classes compatible real-valued solutions, or one should redefine the goal of the problem along the lines of finding a rational-valued approximation to a potentially irrational solution.

As an initial step, we first show that determining whether $r_i < 1$ for a specific bank *i* is at least as hard as solving the square root sum (SQRT-SUM) problem. An instance of SQRT-SUM consists of n + 1 integers $d_1, d_2, ..., d_n, k$ and asks whether $\sum_{i=1}^n \sqrt{d_i} \leq k$. It is known that SQRT-SUM is solvable in PSPACE but it is unknown whether it is in P, or even in NP. In [30] it is shown that it can be solved in polynomial time in the unit-cost RAM model [2, 26, 30].

Lemma 1. For a given pair $(\mathcal{F}, \mathcal{P})$, deciding whether a specific bank is in default is SQRT-SUM-hard.

Proof. We prove the lemma by reducing from SQRT-SUM to CDS-PRIORITY-CLEARING. Let $(d_1, ..., d_n, k)$ be an instance of SQRT-SUM. Firstly, we note that in [4] it is shown that checking whether $\sum_{i=1}^n \sqrt{d_i} = k$ can be done in polynomial time. We check whether equality holds for our input first and proceed without loss of generality to the proof without minding equality.

 $^{^{5}}$ We will make use of these facts in the proof of Theorem 1, below.

We construct a financial system \mathcal{F} as follows, first we construct n financial subnetworks which we refer to as square root gadgets and denote by $g_{i,\sqrt{.}}$ the *i*th square root gadget. Whenever referring to a node k, belonging in a square root gadget $g_{i,\sqrt{.}}$, we use the notation k_i . The square root gadget $g_{i,\sqrt{.}}$ consists of the financial network we presented in Fig. 1, augmented with two additional nodes x_i, y_i , and the CDS contract $(x_i, y_i, 2_i)$. We let the external assets of nodes 2_i and 5_i be $e_{2_i} = e_{5_i} = 1 - d_i$. We let $e_{x_i} = 1$ and $e_{y_i} = 0$ and the CDS contract $(x_i, y_i, 2_i)$ has a notional of 1: $c_{x_i, y_i}^{2_i} = 1$. The n square root gadgets are all connected to a single node τ with $e_{\tau} = 0$, by n debt contracts $\{(y_i, \tau) : i \in [n]\}$. There is one further node τ' to which τ is connected through debt contract (τ, τ') with notional $c_{\tau,\tau'} = k$. The construction is illustrated in Fig. 2.

We claim that this resulting financial system has a clearing recovery rate vector r with $r_{\tau} = 1$ if and only if $\sum_{i=1}^{n} \sqrt{d_i} \ge k$. From the analysis of Example 1, it follows that under any clearing recovery rate vector r, the recovery rate of node 2_i is $r_{2_i} = 1 - \sqrt{d_i}$ for all $i \in [n]$. Since node 2_i is always in default(assuming all $d_i \neq 0$) the CDS $(x_i, y_i, 2_i)$ is activated and since node x_i can fully pay its liabilities, node y_i receives a payment of $1 - r_{2_i} = \sqrt{d_i}$. This implies that τ receives a total payment of $\sum_{i \in [n]} \sqrt{d_i}$. Since τ has only one liability of k, it holds that $r_{\tau} = 1$ if and only if the total payment that τ receives exceeds k, i.e., if and only if $\sum_{i \in [n]} \sqrt{d_i} \ge k$, and this proves the claim.



Fig. 2. The financial system constructed from a given Square Root Sum instance

Next we show that CDS-PRIORITY-CLEARING and its strong approximation variant are FIXP and $FIXP_a$ complete, respectively. Our hardness reduction does not start from a particular FIXP-hard problem. Rather, we show that we can take an arbitrary algebraic circuit and encode it in a direct way in the form of a financial system accompanied by a specific singleton liability profile. Hence, our polynomial time hardness reduction works from any arbitrary fixed point problem in FIXP. The reduction is constructed by devising various financial network gadgets which enforce that certain banks in the system have recovery rates that are the result of applying one of the operators in FIXP's arithmetic base to the recovery rates of two other banks in the system: In other words, we can design our financial systems and singleton liability priorities such that the interrelation between the recovery rates mimics a computation through an arbitrary algebraic circuit.

Theorem 1. CDS-PRIORITY-CLEARING is FIXP-complete, and its strong approximation version is $FIXP_a$ -complete.

Proof (sketch). Containment of CDS-PRIORITY-CLEARING in FIXP is immediate: The clearing vectors for an instance $I = (N, e, c) \in \text{CDS-PRIORITY-CLEARING}$ are the fixed points of the function f_I defined coordinate-wise by $f_I(r)_i = \frac{a_i(r)}{\max\{l_i(r), a_i(r)\}}$ as in (6). The functions $l_i(r)$ and $a_i(r)$ are defined in (1) and (4), from which it is clear that f_I can be computed using a polynomial size algebraic circuit with only $\{\max, +, *\}$, and rational constants. Note that non-degeneracy of I prevents division by 0, so that the output of the circuit is well-defined for every $x \in [0, 1]^n$. This shows that CDS-PRIORITY-CLEARING is in FIXP and that its strong approximation version is in FIXP_a.

For the FIXP-hardness of the problem, let Π be an arbitrary problem in FIXP. We describe a polynomial-time reduction from Π to CDS-PRIORITY-CLEARING. Let $I \in \Pi$ be an instance, let $F_I : [0,1]^n \to [0,1]^n$ be I's associated fixed point function, and let C_I be the algebraic circuit corresponding to F_I . Our reduction is analogous to the proof in [17], where the variant of the problem with proportional payments is shown to be FIXP-hard: The proof consists of a "pre-processing" step (in which the algebraic circuit is transformed into a specific form) followed by a transformation into a financial network, where a set of financial subnetwork gadgets are interconnected, and each such gadget corresponds to an arithmetic gate in the algebraic circuit. The pre-processing step of this reduction is entirely equal to that of [17]: This step transforms C_I into a circuit C'_I that satisfies that all the signals propagated by all gates in C'_I and all the used rational constants in C'_I are contained in the interval [0, 1]. The transformed circuit C'_I may contain two additional type of gates: Division gates and gates that compute the absolute value of the difference of two operands. We will refer to the latter type of gate as an absolute difference gate. The circuit C'_I does not contain any subtraction gates, and will not contain max and min gates either. The reader interested in the details of this pre-processing step is referred to [17].

The second part of the proof (the transformation step) differs from [17] in that a different set of gadgets needs to be used. For notational convenience, we may treat C'_I as the function F_I , hence we may write $C'_I(x) = y$ to denote $F_I(x) = y$.

Each of our gadgets has one or two *input banks* that correspond to the input signals of one of the types of arithmetic gate, and there is an *output bank* that corresponds to the output signal of the gate. For each of the gadgets, it holds that the output bank must have a recovery rate that equals the result of applying the respective arithmetic operation on the recovery rates of the input banks. Each gadget represents an arithmetic gate of the circuit, and the output banks of a gadget are connected to input banks of other gadgets such that there is a direct correspondence with the infrastructure of the arithmetic circuit. For precise details on this correspondence, we refer the reader to the proof in [17].

For defining our gadgets, we use our graphical notation for convenience. As stated, our gadgets each have one or two input banks, and one output bank, where the output bank is forced to have a clearing recovery rate that is obtained by applying a arithmetic operation on the recovery rate of the input banks, so as to simulate the arithmetic basis on which the algebraic circuit C'_I is built.

As a simple example of one of these gadgets we define a straightforward addition gadget, named g_+ , depicted in Fig. 3a. This gadget directly accounts for the addition operation in the arithmetic basis. In our figures, input banks are denoted by black arrows incoming from the left, and output banks correspond to black arrows pointing out of the bank. The black arrows represent connections to other gadgets, and these connections are always realised by a debt contract of unit cost, and are always from the output node of a gadget to an input node of another gadget. Another slightly more complex gadget example is the *positive* subtraction gadget, displayed in Fig. 3b, taking two inputs r_1, r_2 and producing as output the value max $\{0, r_1 - r_2\}$. This gadget is in turn used to form the absolute difference gadget, which can be constructed by combining two positive subtraction gadgets with an addition gadget (as $|r_1 - r_2| = \max\{0, r_1 - r_2\} + \max\{0, r_2 - r_1\}$). It is also used in the construction of our multiplication gadget.

In the figures representing our gadgets, some of the nodes have been annotated with a formula in terms of the recovery rates of the input banks of the gadget, subject to the resulting values being in the interval [0, 1]. This can be seen for example in the output node of our addition gadget in Fig. 3a. Such a formula represents the value that a clearing recovery rate must satisfy for the respective node. It is straightforward to verify that the given formulas are correct for each of our gadgets.

Since all signals inside C'_I are guaranteed to be in [0, 1] for all input vectors, our financial system gadgets can readily be used and connected to each other to construct a financial system that corresponds to C'_I , i.e., such that the clearing recovery rates of the input and output banks of each of the gadgets must correspond to each of the signals inside the circuit C'_I .

The remaining gadgets are omitted. Some of our gadgets are composed of auxiliary gadgets, where in particular, the construction of the division gadget is rather involved, and a separate discussion of how this gadget is constructed is deferred to the full version of the paper. Another notably non-trivial gadget is the multiplication gadget, where the main challenge in its construction is to ensure that non-degeneracy holds and that the expressions for the intermediate recovery rates inside the gadget are all contained in [0, 1].

4 Hardness of Deciding the Best Priority Profile

In a financial network, each node i can be assigned one of outdeg(i)! singleton liability priority lists. Consequently the number of candidate priority profiles for a system is exponentially large in terms of its input size. Next, we consider the setting where a financial authority is able to determine which priority list each bank should get assigned, and is interested in assigning these in such a way that



(a) Addition gadget g_+ .



(b) Positive subtraction gadget $g_{\text{pos}-}$. Generates the value max $\{0, r_1 - r_2\}$ from the input recovery rates r_1 and r_2 . In this gadget node 6 pays according to the priority list $P_6 = ((6,7,4) \mid (6,8))$ thus it submits to node 8 the amount of $r_1 - r_2$ in case $r_1 > r_2$ and 0 otherwise

Fig. 3. Example gadgets from our reduction.

a specific objective is optimised for. We show that this problem is NP-hard for a set of natural choices of objective functions:

- 1. Minimising the number of defaulting nodes.
- 2. Minimising the number of not fully paid liabilities.
- 3. Minimising the number of activated CDSes in the financial system.
- 4. Maximising the equity of a specific node.
- 5. Maximising the liquidity in the financial system.

We here only prove Statements 1, 2 and 4.

Lemma 2. Finding a priority list profile that minimises the number of defaulting banks and finding a priority list profile that minimises the number of not fully paid liabilities are both NP-hard problems.

Proof. We prove the lemma via a reduction from the satisfiability problem (SAT), where we are given a boolean formula in conjunctive normal form, and have to determine whether there is a truth assignment to the variables that renders the formula true. Let $F = \bigwedge_{i=1}^{n} C_i$ be a SAT instance, where C_1, \ldots, C_n are the clauses, and let V_F be the set of all variables that in F. We create a financial network from F as follows. For each variable $x \in V_F$ we construct a gadget that is refereed to as the x-subnetwork. Each x-subnetwork consists of four nodes, labeled as $i_x, x, \neg x, j_x$, where $e_{i_x} = 1$ and $e_x = e_{\neg x} = e_{j_x} = 0$, and of four debt contracts, $\mathcal{DC} = \{(i_x, x), (i_x, \neg x), (x, j_x), (\neg x, j_x)\}$ all with contract notionals equal to one. Moreover, the constructed financial network has n further nodes, labeled as C_1, \ldots, C_n , that correspond to each clause in F with $e_{C_1} = \cdots = e_{C_n} = 0$, and one terminal node labeled as τ , towards which each C_i

holds a debt contract of notional one, i.e. $c_{C_i,\tau} = 1$. Finally for each variable x of F, we construct two nodes labeled as k_x and $k_{\neg x}$ respectively, where e_{k_x} and $e_{k_{\neg x}}$ are equal to the number of occurrences of literals x and $\neg x$ in F, respectively. Whenever a literal l belongs to a clause C_i we construct the CDS $(k_l, C_i, \neg l)$ with $c_{k_l,C_i}^{-1} = 1$. An example of a network induced from $F = (x \lor y) \land (y \lor \neg y)$ is given in Fig. 4. We now map a truth assignment $T : V_F \mapsto \{\text{true}, \text{false}\}$ to a priority list profile \mathcal{P}_T as follows: If T(x) = true, node i_x prioritises node x, and otherwise it prioritises node $\neg x$. All k_l nodes posses enough external assets to fully pay their debts under any priority list, so we can take an arbitrary list for those nodes, and all remaining nodes have at most one liability. Conversely, from a priority list profile \mathcal{P} we induce the truth assignment $T_{\mathcal{P}}$ as follows: if i_x prioritises x then T(x) = true otherwise T(x) = false.

Let $T: V_F \mapsto \{\text{true, false}\}\$ be any truth assignment and consider the priority list profile \mathcal{P}_{T} . In each x-subnetwork, node i_x can fully pay only its first priority, thus in each x-subnetwork there exist two defaulting nodes regardless of the choice of priority list of i_x , meaning that the minimum number of defaulting nodes in the induced financial system is $2|V_F|$. Similarly, each x-subnetwork has two not fully paid liabilities regardless of the choice of priority lost of i_x . The only additional defaulting nodes might be the nodes C_1, \ldots, C_n , and the only additional not fully paid liabilities might be on the n debt contracts in which one of C_1, \ldots, C_n is the debtor. Let us inspect which of the latter set of nodes are defaulting and which of the latter liabilities are not fully paid under \mathcal{P}_{T} . If clause C_i is a clause that is not satisfied under T, then none of the CDSes involving node C_i are activated, and node C_i does not have any assets under \mathcal{P}_{T} . Since C_i has to pay 1 to τ , node C_i will be in default, and C_i 's liability of 1 will not be paid. If clause C_i is a clause that is satisfied under T, then at least one CDS involving node C_i is activated, and since the reference bank in this CDS has recovery rate 0, node C_i will receive the CDS's full notional of 1, with which it can fully pay its liability of 1. Hence, in the latter case, node C_i is not in default.

Hence, for a truth assignment T, under \mathcal{P}_{T} , the number of banks in default and the number of not fully paid liabilities is are both equal to $2|V_F|$ plus the number of unsatisfied clauses. Since we argued above that restricting to the profiles { $\mathcal{P}_{T} \mid T$ is a truth-assignment for F} is without loss of generality, from finding the profile of priority lists minimising the number of defaulting banks or minimising the number of not fully paid liabilities in the constructed financial network, one can infer whether the formula F is satisfiable, which proves our claim.

Lemma 3. Finding a priority list profile that maximises the equity of a specific node is NP-hard.

Proof. We prove the lemma via a reduction from KNAPSACK. Let us be given a knapsack instance with a knapsack of capacity B, a set $S = \{a_1, \ldots, a_n\}$ of objects having profit profit (a_i) and size size (a_i) . Without loss of generality, we assume that size (a_i) , for all $a_i \in S$ as well as B are integer numbers. We



Fig. 4. The financial system corresponding to the formula $F = \{x \lor \neg y\} \land \{y \lor \neg y\}$.

construct a financial network $\mathcal{F} = (N, e, c)$ as follows: We introduce a node 0 with external assets $e_0 = B$ and for each object $a_j \in S$ we introduce a corresponding node j and let node 0 hold a debt contract towards each node j with notional $c_{0,j} = \operatorname{size}(a_j)$. Moreover we introduce a node τ with $e_{\tau} = 0$, and a node Twith $e_T = 0$, which we refer to as the *terminal node*. We add a debt contract of notional $\operatorname{size}(a_j)$ from each node j to node τ . For each node j we moreover introduce a j-subnetwork, consisting of:

- two nodes y_j and x_j with $e_{y_j} = 1, e_{x_j} = 0$,
- a CDS contract (y_j, x_j, j) with notional $c_{y_j, x_j}^j = \max_{a_j \in S} \{size(a_j)\},\$
- a node z_j towards which x_j holds a debt contract of notional $c_{x_j,z_j} = 1$.
- a node k_j with $e_{k_j} = \text{profit}(a_j)$ and an outgoing CDS (k_j, T, x_j) with notional $\text{profit}(a_j)$.

The construction of the *j*-subnetwork is illustrated in Fig. 5.

Assume an optimal solution to the original Knapsack instance and let OPT be the set of the objects a_j contained in it. We know that $\sum_{a_i \in OPT} \operatorname{size}(a_i) \leq B$ and that $\sum_{a_i \in OPT} \operatorname{profit}(a_i)$, is the maximum profit that can fit in the knapsack. We define the set $N = \{1, \ldots, n\}$, and $N_{OPT} = \{j \mid a_j \in OPT\}$ to be the set containing all nodes in \mathcal{F} that correspond to objects contained in the optimal solution. Fix the singleton liability priority list for node 0 to be $\mathcal{P}_0 = (\{N_{OPT}\} \mid \{N \setminus N_{OPT}\})$, meaning that node 0 first prioritises all creditors in N_{OPT} in an arbitrary order and afterwards all other creditors in $N \setminus N_{OPT}$ again in an arbitrary order. Next we prove that under this profile, node Treceives its maximum total assets. Observe that $\forall j \in N_{OPT}, p_{0j} = \operatorname{size}(a_j)$ since $\sum_{j \in N_{OPT}} c_{0j} = \sum_{a_j \in OPT} \operatorname{size}(a_j) \leq B = e_0$. Since every j node that corresponds to an object $a_j \in OPT$ receives size (a_j) , it can fully pay node τ , so $r_j = 1$. For all creditors $m \in N \setminus N_{OPT}$ it holds that $p_{0m} < \operatorname{size}(a_m)$



Fig. 5. The j subnetwork of the constructed financial system

So, $\forall j \in N_{\text{OPT}} : r_j = 1$ while $\forall m \in N \setminus N_{\text{OPT}} : r_m < 1$. Next we prove that for each $j \in N_{\text{OPT}}$, $p_{k_j,T} = \text{profit}(a_j)$ and for each $m \in N \setminus N_{\text{OPT}}$, $p_{k_m,T} = 0$. Take a $j \in N_{\text{OPT}}$, we know that $r_j = 1$, which implies that the CDS (y_j, x_j, j) is not activated thus $r_{x_j} = 0$ which in turn activates the CDS (k_j, T, x_j) where node k_j pays $\text{profit}(a_j)$ to node T. On the other side, for a $m \in N \setminus N_{\text{OPT}}$, it holds that $r_m < 1$, which means that the CDS (y_m, x_m, m) is activated and generates a liability of $\max_{a_i \in S} \{\text{size}(a_i)\} \cdot (1 - r_m)$ for node y_m . We prove that this liability is at least 1: For an object $a_m \notin \text{OPT}$, r_m indicates the proportion of $\text{size}(a_m)$ that fits in the available knapsack area unoccupied by the objects in OPT. Obviously for $a_m \notin \text{OPT}$, $\text{size}(a_m) > B - \text{size}(\text{OPT})$, otherwise $a_m \in \text{OPT}$ and $r_m \cdot \text{size}(a_m) + \text{size}(\text{OPT}) = B$. Since by assumption B and $\text{size}(a_j)$ for all $a_j \in S$ are integers, it holds that $\forall a_m \notin \text{OPT}$, $r_m \leq (\max_{a_k \in S} \{\text{size}(a_k)\} - 1)/(\max_{a_k \in S} \{\text{size}(a_k)\}$, so the liability for y_m is:

$$l_{y_m, x_m}^m = \max_{a_k \in S} \{ \text{size}(a_k) \} (1 - r_m) \\ \ge \max_{a_k \in S} \{ \text{size}(a_k) \} \left(1 - \frac{\max_{a_k \in S} \{ \text{size}(a_k) \} - 1}{\max_{a_k \in S} \{ \text{size}(a_k) \}} \right) = 1.$$

So eventually, $\forall m \in N \setminus N_{\text{OPT}}$, $p_{y_m,x_m} = 1$. Now $r_{x_m} = 1$ thus the CDS (k_m, T, x_m) is not activated meaning that $p_{k_m,T} = 0$. From the above observations we conclude that the equity of T is $\sum_{j \in N_{\text{OPT}}} \operatorname{profit}(a_j)$: node T receives money from all nodes that correspond to objects contained in OPT. We claim that this is the maximum equity node T can receive. If there exist a higher equity for T, then this must be generated from another profile \mathcal{P}'_0 that corresponds to a solution to the original Knapsack instance with higher profit than the optimal one which is a contradiction.

For the opposite direction assume \mathcal{P}_0 to be the profile of node 0 that maximises T's equity. Let $A = \{a_j \mid p_{0j} = \operatorname{size}(a_j)\}$ be the set of objects that corresponds to creditor nodes that node 0 can fully pay. Obviously Acan be computed in polynomial time from \mathcal{P}_0 . We claim that A is the optimal solution to the original Knapsack instance. Assume that there exists another set A' such that $\sum_{a_j \in A'} \operatorname{size}(a_j) \leq B$ and $\sum_{a_j \in A'} \operatorname{profit}(a_j) >$ $\sum_{a_j \in A} \operatorname{profit}(a_j)$. Now node 0 could rearrange its priorities by prioritising all creditors j for which $a_j \in A'$. Doing so, node 0 can fully pay all nodes jfor which $a_j \in A'$ since $\sum_{a_j \in A'} \operatorname{size}(a_j) \leq B = e_0$ and node T receives $\sum_{a_j \in A'} \operatorname{profit}(a_j) > \sum_{a_j \in A} \operatorname{profit}(a_j)$, a contradiction to the original assumption that $\sum_{a_j \in A} \operatorname{profit}(a_j)$ is the maximum equity for node T.

5 Conclusions and Future Work

Financial networks have emerged as a fertile research area in both computational complexity and algorithmic game theory. It is paramount to understand systemic risk in finance from both these perspectives. In this paper, we join both streams of work by settling questions around the computational complexity of systemic risk for priority payments, a scheme so far only studied from the game-theoretic point of view. In an interesting parallel with the state of the art for proportional payments, we prove that computing clearing recovery rates is FIXP-complete whereas it is NP-hard to compute the priority lists optimising several measures of financial health of the system.

Our work paves the way for studying payment schemes in financial networks in more detail. Is there a payment scheme for financial networks with derivatives that makes the computation of systemic risk easy and/or that induces "nice" equilibria? We also wonder the extent to which the flexibility of working with financial networks can lead to a deeper understanding of FIXP; e.g., are there payment schemes that can be proved complete for variants of FIXP defined upon a different basis?

References

- Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A.: Systemic risk and stability in financial networks. Am. Econ. Rev. 105(2), 564–608 (2015)
- Bertoni, A., Mauri, G., Sabadini, N.: A characterization of the class of functions computable in polynomial time on random access machines. In: STOC, pp. 168– 176. ACM (1981)
- Bertschinger, N., Hoefer, M., Schmand, D.: Strategic payments in financial networks. In: ITCS. LIPIcs, vol. 151, pp. 46:1–46:16 (2020)
- Borodin, A., Fagin, R., Hopcroft, J.E., Tompa, M.: Decreasing the nesting depth of expressions involving square roots. J. Symb. Comput. 1(2), 169–188 (1985)
- Cifuentes, R., Ferrucci, G., Shin, H.S.: Liquidity risk and contagion. J. Eur. Econ. Assoc. 3(2–3), 556–566 (2005)
- Eisenberg, L., Noe, T.H.: Systemic risk in financial systems. Manag. Sci. 47(2), 236–249 (2001)
- Elliott, M., Golub, B., Jackson, M.O.: Financial networks and contagion. Am. Econ. Rev. 104(10), 3115–3153 (2014)
- Etessami, K., Yannakakis, M.: On the complexity of nash equilibria and other fixed points. SIAM J. Comput. **39**(6), 2531–2597 (2010)
- Filos-Ratsikas, A., Giannakopoulos, Y., Hollender, A., Lazos, P., Poças, D.: On the complexity of equilibrium computation in first-price auctions. In: EC, pp. 454–476 (2021)
- Filos-Ratsikas, A., Hansen, K.A., Høgh, K., Hollender, A.: Fixp-membership via convex optimization: games, cakes, and markets. In: FOCS, pp. 827–838 (2021)
- Garey, M.R., Graham, R.L., Johnson, D.S.: Some NP-complete geometric problems. In: STOC, pp. 10–22 (1976)
- Glasserman, P., Young, H.P.: How likely is contagion in financial networks? J. Bank. Finance 50, 383–399 (2015)

- Goldberg, P.W., Hollender, A.: The hairy ball problem is PPAD-complete. J. Comput. Syst. Sci. 122, 34–62 (2021)
- Heise, S., Kühn, R.: Derivatives and credit contagion in interconnected networks. Eur. Phys. J. B 85(4), 1–19 (2012)
- Hemenway, B., Khanna, S.: Sensitivity and computational complexity in financial networks. Algorithmic Finance 5(3–4), 95–110 (2016)
- Hu, D., Zhao, J.L., Hua, Z., Wong, M.C.S.: Network-based modeling and analysis of systemic risk in banking systems. MIS Q. 1269–1291 (2012)
- Ioannidis, S.D., De Keijzer, B., Ventre, C.: Strong approximations and irrationality in financial networks with derivatives. In: ICALP (2022)
- Kanellopoulos, P., Kyropoulou, M., Zhou, H.: Financial network games. CoRR, abs/2107.06623 (2021)
- Kanellopoulos, P., Kyropoulou, M., Zhou, H.: Forgiving debt in financial network games. CoRR, abs/2202.10986 (2022)
- Papadimitriou, C.H.: The euclidean traveling salesman problem is np-complete. Theor. Comput. Sci. 4(3), 237–244 (1977)
- Papp, P.A., Wattenhofer, R.: Network-aware strategies in financial systems. In: ICALP, vol. 168, pp. 91:1–91:17 (2020)
- Papp, P.A., Wattenhofer, R.: Debt swapping for risk mitigation in financial networks. In: EC, pp. 765–784 (2021)
- Papp, P.A., Wattenhofer, R.: Default ambiguity: finding the best solution to the clearing problem. In: Feldman, M., Fu, H., Talgam-Cohen, I. (eds.) WINE 2021. LNCS, vol. 13112, pp. 391–409. Springer, Cham (2022). https://doi.org/10.1007/ 978-3-030-94676-0 22
- Papp, P.A., Wattenhofer, R.: Sequential defaulting in financial networks. In: ITCS, vol. 185, pp. 52:1–52:20 (2021)
- Rogers, L.C.G., Veraart, L.A.M.: Failure and rescue in an interbank network. Manag. Sci. 59(4), 882–898 (2013)
- Schönhage, A.: On the power of random access machines. In: Maurer, H.A. (ed.) ICALP 1979. LNCS, vol. 71, pp. 520–529. Springer, Heidelberg (1979). https:// doi.org/10.1007/3-540-09510-1_42
- Schuldenzucker, S., Seuken, S.: Portfolio compression in financial networks: incentives and systemic risk. In: EC, p. 79 (2020)
- Schuldenzucker, S., Seuken, S., Battiston, S.: Clearing payments in financial networks with credit default swaps. In: EC, p. 759 (2016)
- Schuldenzucker, S., Seuken, S., Battiston, S.: Finding clearing payments in financial networks with credit default swaps is PPAD-complete. In: ITCS, vol. 67, pp. 32:1– 32:20 (2017)
- Tiwari, P.: A problem that is easier to solve on the unit-cost algebraic RAM. J. Complex. 8(4), 393–397 (1992)
- Yannakakis, M.: Equilibria, fixed points, and complexity classes. In: STACS, vol. 1, pp. 19–38 (2008)