

Liquidity in Credit Networks: A Little Trust Goes a Long Way

Pranav Dandekar*

Ashish Goel†

Ramesh Govindan‡

Ian Post§

Abstract

Credit networks represent a way of modeling trust between entities in a network. Nodes in the network print their own currency and trust each other for a certain amount of each other’s currency. This allows the network to serve as a decentralized payment infrastructure—arbitrary payments can be routed through the network by passing IOUs between trusting nodes in their respective currencies—and obviates the need for a common currency. Nodes can repeatedly transact with each other and pay for the transaction using trusted currency. A natural question to ask in this setting is: how long can the network sustain liquidity, *i.e.*, how long can the network support the routing of payments before credit dries up? We answer this question in terms of the long term failure probability of transactions for various network topologies and credit values.

We prove that the transaction failure probability is independent of the path along which transactions are routed. We show that under symmetric transaction rates, the transaction failure probability in a number of well-known graph families goes to zero as the size, density or credit capacity of the network increases. We also show via simulations that even networks of small size and credit capacity can route transactions with high probability if they are well-connected. Further, we characterize a centralized currency system as a special type of a star network (one where edges to the root have infinite credit capacity, and transactions occur only between leaf nodes) and compute the steady-state transaction failure probability in a centralized system. We show that liquidity in star networks, complete graphs and Erdős-Rényi networks is comparable to that in equivalent centralized currency systems; thus we do not lose much liquidity in return for their robustness and decentralized properties.

1 Introduction

One of the primary functions of money is as a medium of exchange. This function is facilitated by governments and central banks that issue currency and declare it to be legal tender, *i.e.*, the state promises to accept that currency back as payment for goods and services. Thus currency represents an obligation issued by the state, and when used to make a payment, results in a transfer of the state’s obligation from the payer to the payee. A decision to accept payment in a currency is therefore a decision to trust the issuer of the currency to fulfill its obligations.

The modern banking system is a centralized currency infrastructure. The central bank sits at the root of the tree. It issues currency, government notes, etc. Retail banks, individuals and businesses form the leaves of this tree. Since they cannot print their own currency, they trust the central bank for an infinite amount of money (in theory, at least). Many non-governmental organizations also issue their own currencies, known as *scrip*. Examples of such scrip systems in the online world include Facebook credits, Yootles [RSK07], P2P systems where such currency is used to solve the free-rider problem [VCS03, AJ06], and virtual worlds such as Everquest and Second Life that have their own currencies. A challenge with scrip systems is to determine how much currency to issue in order to ensure liquidity without devaluing the currency [KFH07].

We study an alternate model of credit, credit networks, introduced independently by DeFigueiredo and Barr [DB05], by Ghosh et al. [GMR+07], and by Karlan et. al. [KMRS09]. We generalize the model based on the insight that payments can be made in any currency as long as it is trusted by the payee. So nodes in our model act as banks and print their own currency and trust each other for certain amounts of each other’s currency. This allows them to transact with each other without the need for a common currency: as long as there is a chain of sufficient credit from the payee to the payer, arbitrary payments can be made by passing obligations between trusting agents in their respective currencies. We study the question of long-term liquidity (*i.e.*, capacity to route payments) in credit networks when nodes repeatedly transact with each other.

Illustrative Example. Consider a credit network with three nodes u, v and w with edge (u, v) having credit limit c_1 (*i.e.*, u trusts v for up to c_1 units of v ’s currency) and edge (v, w) having a credit limit c_2 . If w wants to purchase a product or service from u worth p units of w ’s currency, where $p \leq \min\{c_1, c_2\}$, the transaction proceeds by w issuing an IOU to v worth p units of w ’s currency and v issuing an IOU to u worth p units of v ’s currency. However, if $p > \min\{c_1, c_2\}$ the transaction fails. As a

*Department of Management Science & Engineering, Stanford University. Email: ppd@stanford.edu.

†Departments of Management Science & Engineering and (by courtesy) Computer Science, Stanford University. Email: ashishg@stanford.edu.

‡Computer Science Department, University of Southern California. Email: ramesh@usc.edu.

§Computer Science Department, Stanford University. Email: itp@stanford.edu.

result of a successful transaction, the edges (u, v) and (v, w) now have capacities $c_1 - p$ and $c_2 - p$ respectively (since v and w have depleted part of their credit line). But this transaction also results in the creation of edges (v, u) and (w, v) each having capacity p , since now v can trust u for up to p units of its own currency, and w can trust v for up to p units of its own currency. Thus routing payments in credit networks is identical to routing residual flows in general flow networks.

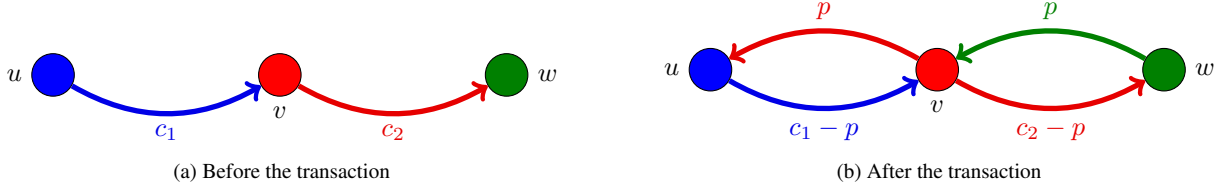


Figure 1: Illustrative Example

The papers that introduced this model used it as a generalization of budget constraints in auctions [GMR⁺07] and as a mechanism for providing social collateral/insurance for transactions between untrusting agents in a network [DB05, KMRS09]. In addition to these papers, Resnick and Sami [RS09] study how to update credit values along edges over time while ensuring robustness to Sybil attacks. They show that there is a social cost to allowing transactions via intermediaries—any protocol for updating credit values that is Sybil-proof in a certain sense can lead to a reduction in overall credit balances in expectation. Previous work on this model used the term “trust network” for it; we prefer the term “credit network” to distinguish this model from all the previous ones that also model trust/reputation in networks in settings unrelated to ours such as web search [PBMW99, GGMP04] and studying people’s opinions of each other [GKRT04, LHK10].

Credit networks have a number of interesting and useful properties. First, the loss incurred by nodes in the network is *bounded*. Consider a scenario where a potential attacker, say node x , injects nodes into this network. The attacker’s nodes can extend each other credit lines worth arbitrary amount of money, but in order to transact with one of the honest players the attacker x must dupe one of them, say u , into extending him a line of credit. If u extends x a credit line for c units, x can transact with honest nodes in the network for up to c units by giving u IOUs worth c units in its currency. Once the transaction is complete, x can then disappear leaving u with unredeemable IOUs worth c units, and u incurs a loss equal to c units in x ’s currency. This demonstrates that the total loss incurred by the nodes in the network is bounded by the total credit extended by honest nodes to the attacker’s nodes. In particular, the loss incurred by the nodes in the network does not depend on the number of nodes the attacker injects. Second, since a node in the network only accepts IOUs from other nodes to whom it has extended a line of credit, the loss incurred by honest players in the network is *localized*. In other words, the only nodes that incur a loss are the ones that extended credit to the attacker nodes. In this sense, the system is fair; it protects a node from losses resulting from another node’s lack of judgement. Finally, routing payments in credit networks is efficient since it only requires a max-flow computation [DB05, GMR⁺07].

These properties make this model useful not only for monetary transactions, but also in any setting where there is a need to model trust between nodes in a network. It is particularly well-suited for transactions in exchange economies such as P2P networks where it can be used to improve inefficiencies resulting from asynchronous demand and bilateral trading [LHL⁺10]. It has been used as part of a system called Ostra [MPGD08] designed to thwart spam in email. Ostra is similar in spirit to Gmail’s Priority Inbox. This model can also be used in settings such as combating spam in viral marketing over social networks and packet routing in mobile ad-hoc networks. There is a large body of work in economics and sociology on social capital and favor exchanges in networks [JRBT11]. This model provides a rigorous way of not only keeping track of favors owed to and by each individual in a network, but also a way to exchange favors, via trusted intermediaries, between individuals that do not know each other directly.

However, in order for the model to be of practical use it should be able to support repeated transactions between nodes over a long period of time. This motivates the following question, which we formulate and study in this paper: if the network is sufficiently well-connected and has sufficient credit to begin with, can we sustain transactions in perpetuity without additional injection of credit? How does liquidity depend upon network topology and transaction rates between nodes, and how does it compare with the centralized model described above?

We show positive results for all of the above questions. We show that the failure probability of transactions is independent of the path used to route flow between nodes. For symmetric transaction rates, we show analytically and via simulations that the failure probability for complete graphs, Erdős-Rényi graphs, and preferential attachment graphs goes to zero as the size, the density or the credit capacity of the network increases. Further, we characterize a centralized currency system as a special type of a star network (one where edges to the root have infinite credit capacity, and transactions occur only between leaf nodes) and compute its steady-state failure probability. We show how to construct a centralized system that is equivalent to a given

credit network and that the steady-state failure probability in star networks and complete graphs is within a constant-factor of that in equivalent centralized currency systems. We show that the node bankruptcy probability in credit networks is $\Theta(1/h_G)$ whereas that in an equivalent centralized system is $\Theta(1/\bar{c}_G)$, where h_G is the harmonic mean of the total credit capacities on edges incident upon nodes in G and \bar{c}_G is the corresponding arithmetic mean. These results are similar in spirit to price of anarchy bounds—in return for all the benefits resulting from the decentralized nature of the system, we do not lose much liquidity compared to a centralized model with a common currency.

We would like to qualify our results by stating that nodes in our model are not endowed with any rationality. We treat credit values on edges and transaction rates between nodes as being exogenously defined. We also ignore currency exchange ratios, *i.e.*, we assume them to be unity for clarity of exposition; this does not qualitatively affect our results. These questions represent promising directions for further study of this model and we discuss them at the end of the paper along with other open problems.

The rest of the paper is organized as follows: Section 2 defines our model and describes our main results, Section 3 contains steady-state analysis of the transaction success probability in our model and its comparison with a centralized model of credit, Section 4 discusses results of simulating repeated transactions on two random graph families, Erdős-Rényi graphs and power-law graphs constructed using the Barabási-Albert preferential attachment model [BA99], and finally Section 5 has some concluding remarks along with a discussion of a number of interesting directions for further study of this model.

2 Our Model & Results

A credit network $G = (V, E)$ is a directed graph with n nodes and m edges. Nodes represent entities or agents. Edges represent pairwise credit limits between agents. An edge $(u, v) \in E$ has capacity $c_{uv} > 0$, which implies that u has extended a credit line of c_{uv} units to v in v 's currency. If a node s needs to pay p units in her currency to node t (for example, to buy a good that t is selling), the payment can go through if the maximum credit flow from t to s is at least p units. The payment will get routed through a chain of trusted nodes from s to t . Note that payment flows in the opposite direction of credit, so a payment merely results in a redistribution of credit; the total credit in the network remains unchanged.

We study liquidity under the following simple model of repeated transactions: we assume that each edge (u, v) in G has total credit capacity $c_{uv} + c_{vu} = c$ where c_{uv}, c_{vu} are integers. The transaction rate between nodes is given by the $n \times n$ matrix $\Lambda = \{\lambda_{uv} : u, v \in V; \lambda_{uu} = 0\}$ such that $\sum_{u,v} \lambda_{uv} = 1$. At each time step, we pick a pair of nodes (s, t) with probability λ_{st} and try to route a unit payment along the shortest feasible path from s to t . If such a path exists, we route the flow along that path and modify edge capacities along the path as described previously (*i.e.*, for each edge (u, v) on the path, increment c_{uv} by 1 and decrement c_{vu} by 1). The transaction fails if there is no path between s and t .

We begin by noting that at any time step an edge between nodes u and v can be in one of $(c + 1)$ states, having capacity c from u to v and 0 from v to u , having capacity $c - 1$ from u to v and 1 from v to u and so on down to having 0 capacity from u to v and c from v to u . Each successful transaction results in a change of state for all edges along the shortest path from the source to the destination. Therefore, these repeated transactions between nodes in G induce the following Markov chain $\mathcal{M}(G, \Lambda)$: each state of the Markov chain encapsulates the state of all the edges in the network. Since there are m edges, each of which can be in one of $(c + 1)$ states, the Markov chain has $(c + 1)^m$ states. A successful transaction results in a transition between two states. Let P be the transition matrix. The transition probability, $P(\mathcal{S}_1, \mathcal{S}_2)$ from state \mathcal{S}_1 to \mathcal{S}_2 is equal to λ_{st} where routing a unit payment from s to t in state \mathcal{S}_1 along the shortest feasible path results in the network being in state \mathcal{S}_2 . The probability $P(\mathcal{S}, \mathcal{S})$ is the probability that the transaction is infeasible when the network is in state \mathcal{S} . This is equal to the sum of all probabilities λ_{st} such that the network in state \mathcal{S} has no path from s to t .

A note on terminology: we will sometimes refer to an edge (u, v) as being “bidirectional” if both c_{uv} and c_{vu} are non-zero, and as “unidirectional” if exactly one of c_{uv} and c_{vu} is zero.

We describe our main results below. Our first result obviates the need to route payments along the shortest feasible path.

Theorem. *Let $(s_1, t_1), (s_2, t_2), \dots, (s_T, t_T)$ be the set of transactions of value v_1, v_2, \dots, v_T respectively that succeed when the payment from s_i to t_i is routed along a path \mathcal{P}_i . Then the same set of transactions succeed when the payment from s_i to t_i is routed along any other feasible path \mathcal{P}'_i .*

In order to prove this, we observe that routing flow along a directed cycle changes the state of the network but does not affect the total credit available to any node. This motivates the following definition:

Definition 1. *Given a credit network G , let \mathcal{S} and \mathcal{S}' be two states of G . We say that \mathcal{S}' is cycle-reachable from \mathcal{S} if the network can be transformed from state \mathcal{S} to state \mathcal{S}' by routing a sequence of payments along feasible cycles (*i.e.*, from a node to itself along a feasible path).*

Network Topology	Transaction Regime	Credit Network	Centralized System
Star	uniform	$\Theta(1/c)$	$\Theta(1/c)$
Line	uniform	$1 - \Theta(c/n^2)$	$\Theta(1/c)$
Cycle	uniform	$1 - \Theta(c/n)$	$\Theta(1/c)$
Complete	symmetric	$\Theta(1/nc)$	$\Theta(1/nc)$
Erdős-Rényi ² : $G_c(n, p)$	symmetric	$\Theta(1/npc)$	$\Theta(1/npc)$
Barabási-Albert ³ : $G_c^{BA}(n, d)$	symmetric	$\Theta(1/dc)$	$\Theta(1/dc)$

Table 1: Failure probability comparison between credit networks and equivalent centralized currency systems

Cycle-reachability defines a partition \mathcal{C} over the set of possible states of the network; each equivalence class in \mathcal{C} is the set of states that are cycle-reachable from each other. We show the following result about the steady-state distribution of \mathcal{M} over equivalence classes in \mathcal{C} :

Theorem. Consider a Markov chain $\mathcal{M}_{\mathcal{S}_0}(G, \Lambda)$ starting in state \mathcal{S}_0 induced by a symmetric transaction rate matrix Λ over nodes in G . Let $\mathcal{C}_{\mathcal{S}_0} \subseteq \mathcal{C}$ be the set of equivalence classes accessible from \mathcal{S}_0 under the regime defined by Λ . Then $\mathcal{M}_{\mathcal{S}_0}$ has a uniform steady-state distribution over $\mathcal{C}_{\mathcal{S}_0}$.

We use this result to bound the steady-state node bankruptcy probability in general graphs and transaction failure probability for trees, cycles and complete graphs. A node being bankrupt is a sufficient but not necessary condition for a transaction with that node as the payer to fail; the necessary and sufficient condition for a transaction to fail is that the network has a directed cut of zero capacity between the payer and the payee. Let d_v be the total credit capacity of all edges incident upon node v in a graph G . Let h_G be the harmonic mean of the d_v s and \bar{d}_G be the arithmetic mean. We show that if $\mathcal{M}(G, \Lambda)$ is ergodic and Λ is symmetric, the average steady-state bankruptcy probability of a node in a credit network is $\Theta(1/h_G)$ whereas that in an equivalent centralized system is $\Theta(1/\bar{d}_G)$. We show that if Λ is uniform, the steady-state transaction success probability for line graphs is $\Theta(c/n^2)$ and that for cycles is $\Theta(c/n)$. We also show that if Λ is symmetric, the steady-state transaction failure probability for complete graphs is $\Theta(1/nc)$. Our results imply that for “thin” graphs such as lines and cycles, the steady-state failure probability goes to one with the size of the network, whereas for well-connected graphs such as star networks, complete graphs and Erdős-Rényi networks¹, the steady-state failure probability goes to zero with the size, the average node degree or the credit capacity in the network.

Further, we characterize a centralized currency system as a special type of a tree network (one where edges to the root have infinite credit capacity, and transactions occur only between leaf nodes) and compute its steady-state failure probability. We show how to construct a centralized system that is equivalent to a given credit network and that the steady-state failure probability in star networks and complete graphs is within a constant-factor of that in equivalent centralized currency systems.

Table 1 shows our main analytical results on liquidity for a number of network topologies and their comparison with equivalent centralized currency systems.

In addition to these analytical results, we simulated repeated transactions on two random graph families: Erdős-Rényi graphs and power-law graphs generated using the Barabási-Albert preferential attachment (BA) model [BA99]. The simulations showed that relatively small values of density and credit capacity were sufficient to attain a transaction success probability in excess of 0.9. We also demonstrate that for both topologies, if we hold the average node degree and credit capacity on edges constant, the size of the network (*i.e.*, number of nodes) had no effect on the steady-state success probability.

3 Analysis

We first analyze the combinatorial structure induced by this model of making payments through passing IOUs. An understanding of this combinatorial structure allows us to characterize the steady-state behavior of the Markov chain induced by repeated transactions for a number of network topologies. Finally, we compare the steady-state success probability in credit networks with various topologies to that in an equivalent centralized currency infrastructure.

¹For a connected Erdős-Rényi network, $G_c(n, p)$, the node bankruptcy probability is $\Theta(1/npc)$.

²We conjecture based on simulations and heuristic calculations that the transaction failure probability is $\Theta(1/npc)$.

³We conjecture based on simulations that the failure probability is $\Theta(1/dc)$.

3.1 Combinatorial Structure

A state of the network can be associated with a *generalized score vector*, which characterizes the credit available to each node in that state.

Definition 2. Given a labeled, directed graph D over n nodes with edges having associated capacities, a vector $V = \langle v_1, \dots, v_n \rangle \in \mathbb{R}_+^n$ is the generalized score vector of D if the total capacity on outgoing edges from i in D is v_i for $i = 1, \dots, n$.

This is a generalization of the score vector defined by the outdegree of nodes in an orientation of a labeled multigraph [KW81]. Note that two cycle-reachable states have the same generalized score vector, and in fact, generalized score vectors characterize classes of cycle-reachable states:

Lemma 1 (Generalization of Proposition 4.10 in [Gio07]). Given a credit network G , two states S and S' of G are cycle-reachable if and only if they have the same generalized score vector.

Moreover, two cycle-reachable states have exactly the same set of feasible transactions:

Lemma 2. For any equivalence class $C \in \mathcal{C}$ of a given network, if a transaction (s, t, v) (i.e., routing a payment of v units from node s to node t) is feasible in some state $S \in C$, it is feasible in all states $S' \in C$.

Using the definition of cycle-reachability and the above observations, we show the following:

Theorem 3. Let $(s_1, t_1), (s_2, t_2), \dots, (s_T, t_T)$ be the set of transactions of value v_1, v_2, \dots, v_T respectively that succeed when the payment from s_i to t_i is routed along a path \mathcal{P}_i . Then the same set of transactions succeed when the payment from s_i to t_i is routed along any other feasible path \mathcal{P}'_i .

Proof. Note that a failed transaction does not change the state of the network, so without loss of generality, we assume that the set of successful transactions (s_i, t_i) occurred in successive timesteps $i = 1$ to T . We prove the result by induction on T . The statement clearly holds for $T = 1$.

Assume that the statement holds for $T = k$. Let the initial state of the network be \mathcal{S}_0 . Let \mathcal{S}_k be the state of the network when, starting from \mathcal{S}_0 , transactions (s_i, t_i) for $i = 1, \dots, k$ were routed along path \mathcal{P}_i . Let \mathcal{S}'_k be the corresponding state of the network after, again starting from \mathcal{S}_0 , the same set of transactions were routed along paths $\mathcal{P}'_i, i = 1, \dots, k$. We will show that \mathcal{S}_k is cycle-reachable from \mathcal{S}'_k .

A successful transaction (s, t) only changes the credit extended to s and t ; the credit extended to any of the intermediate nodes remains unchanged. Since \mathcal{S}_k and \mathcal{S}'_k are the resulting states of the network after the same sequence of payments were successfully routed starting at the same state, the credit distribution (or equivalently, the generalized score vectors) in \mathcal{S}_k and \mathcal{S}'_k must be identical. Therefore, \mathcal{S}_k and \mathcal{S}'_k are cycle-reachable. Thus, from Lemma 2 if the transaction (s_{k+1}, t_{k+1}) is feasible in \mathcal{S}_k then it is also feasible in \mathcal{S}'_k . This proves the result. \square

This result shows that routing in credit networks has a *path-independence property*: the choice of paths along which a sequence of payments are routed does not affect their feasibility. Lemma 1 and Lemma 2 immediately implies the following corollary:

Corollary 4. If a transaction (s, t, v) in some state in equivalence class C_i results in a transition to a state in equivalence class C_j , then the reverse transaction (t, s, v) from any state in C_j will result in a transition to a state in C_i .

Note that the above observations do not require the edge capacities or payment flows to be integral. However, when restricted to a setting where edge capacities are integers and we route unit flows between nodes, the states of the credit network become equivalent to orientations of a labeled multigraph (since an edge of capacity c can be viewed as c labeled edges of unit capacity each). It is known that there is a bijection between score vectors of a labeled multigraph and forests of a labeled multigraph [KW81]. In the rest of the paper, we assume that edge capacities are integers and we route unit flows between nodes. This allows us to use the equivalence between forests, score vectors and cycle-reachable equivalence classes to characterize the steady-state of the Markov chain induced by repeated transactions.

3.2 Steady-state Analysis

Using observations about the combinatorial structure of the model, we next prove that the Markov chain induced by a symmetric transaction rate matrix Λ has a uniform steady-state distribution over \mathcal{C} .

Theorem 5. Consider a Markov chain $\mathcal{M}_{\mathcal{S}_0}(G, \Lambda)$ starting in state \mathcal{S}_0 induced by a symmetric transaction rate matrix Λ over nodes in G . Let $\mathcal{C}_{\mathcal{S}_0} \subseteq \mathcal{C}$ be the set of equivalence classes accessible from \mathcal{S}_0 under the regime defined by Λ . Then $\mathcal{M}_{\mathcal{S}_0}$ has a uniform steady-state distribution over $\mathcal{C}_{\mathcal{S}_0}$.

Proof. As a consequence of the above facts, we can represent transactions as resulting in transitions between equivalence classes in \mathcal{C} instead of transitions between states of $\mathcal{M}(G, \Lambda)$. Let \mathcal{T}_{ij} be the set of transactions (s, t) that result in a transition from some state in C_i to some state in C_j . Then, (overloading the symbol P) we define the transition matrix over equivalence classes in $\mathcal{C}_{\mathcal{S}_0}$ as

$$P(C_i, C_j) = \sum_{(s,t) \in \mathcal{T}_{ij}} \lambda_{st}$$

Note that Corollary 4 implies that $(s, t) \in \mathcal{T}_{ij}$ if and only if $(t, s) \in \mathcal{T}_{ji}$. Further since Λ is symmetric, for any equivalence classes $C_i, C_j \in \mathcal{C}_{\mathcal{S}_0}$, $P(C_i, C_j) = P(C_j, C_i)$. Since P is a symmetric stochastic matrix, the uniform distribution is stationary with respect to P . \square

This immediately gives the following corollary:

Corollary 6. If $\mathcal{M}(G, \Lambda)$ is an ergodic Markov chain induced by a symmetric transaction rate matrix Λ , it has a uniform steady state distribution over \mathcal{C} .

Next we show a sufficient condition for \mathcal{M} to be ergodic.

Lemma 7. If for all $u \neq v$, $\lambda_{uv} > 0$, then \mathcal{M} is ergodic.

Proof. \mathcal{M} has a finite number of states. Since $\forall u, v, u \neq v, \lambda_{uv} > 0$, for all states \mathcal{S} of \mathcal{M} , $P(\mathcal{S}, \mathcal{S}) < 1$. Therefore, \mathcal{M} does not have a sink state. It is irreducible since for any states \mathcal{S} and \mathcal{S}' such that the “edit distance” between \mathcal{S} and \mathcal{S}' is k , the k -step transition probability from \mathcal{S} to \mathcal{S}' , $P^k(\mathcal{S}, \mathcal{S}') > 0$. Finally, to see that \mathcal{M} is aperiodic, consider a state \mathcal{S} such that $P(\mathcal{S}, \mathcal{S}) > 0$. For this state, both $P^2(\mathcal{S}, \mathcal{S})$ and $P^3(\mathcal{S}, \mathcal{S})$ are non-zero. Therefore, state i is aperiodic. Since \mathcal{M} is finite-state, irreducible and aperiodic, it is ergodic. \square

Theorem 5 and Corollary 6 reduce the problem of computing the steady-state transaction failure probability to a counting problem. We instantiate this theorem for various network topologies to characterize their steady-state distribution and infer from it the steady-state failure probability in those topologies. For general graphs, we can use this theorem to bound the bankruptcy probability of a node.

Theorem 8. Let $\mathcal{M}(G, \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in G . Fix a vertex v . The steady-state probability that v goes bankrupt is at most $1/(d_v + 1)$, where $d_v = \sum_{(v,u) \in E} (c_{uv} + c_{vu})$ is total credit capacity of all edges incident upon node v .

Proof. From Corollary 6, we know that $\mathcal{M}(G, \Lambda)$ has a uniform steady-state distribution over the cycle-reachable equivalence classes of G . So the steady-state probability that v goes bankrupt is simply the fraction of equivalence classes in which v is bankrupt.

We know that the total number of cycle-reachable equivalence classes of G is the number of subforests over all n nodes, whereas the number of equivalence classes in which v is bankrupt is equal to the number of subforests on the remaining $n - 1$ nodes [KW81]. Let $F(G)$ be the total number of subforests of G and $F(G_{-v})$ be the number of subforests on nodes other than v . An edge $e = (u, v)$ having total credit capacity c_e can be viewed as c_e labeled edges each of unit capacity. So, each subforest on the remaining $n - 1$ nodes can be extended to a subforest on all n nodes by adding one of the d_v edges incident on v or adding v without adding any edges. Therefore $F(G) \geq (d_v + 1)F(G_{-v})$. Thus, the steady-state probability that v goes bankrupt, $F(G_{-v})/F(G)$, is at most $1/(d_v + 1)$. \square

Therefore, if h_G is the harmonic mean of the d_v s in G , then the average node bankruptcy probability in G is $\Theta(1/h_G)$.

3.2.1 Trees

The above results allow us to make the following observations about the steady-state distribution over trees. We denote by $T_{n;c}$ a tree with n nodes whose each edge has total credit capacity c .

Theorem 9. If $\mathcal{M}(T_{n;c}, \Lambda)$ is an ergodic Markov chain induced by a symmetric transaction rate matrix Λ on a tree $T_{n;c}$, then $\mathcal{M}(T_{n;c}, \Lambda)$ has a uniform steady-state distribution over its states.

Proof. The result follows directly from Corollary 6 and the fact that, since trees have no cycles, each equivalence class in \mathcal{C} is a singleton. \square

Since $\mathcal{M}(T_{n;c}, \Lambda)$ has a uniform steady-state distribution over its states, the steady-state probability that a transaction over a path of length l will succeed is $[c/(c+1)]^l$; this follows from the fact that each edge along the path should have non-zero capacity in the direction of the transaction. Therefore, the steady-state success probability, \mathbb{P}_s , for $T_{n;c}$ is given by

$$\mathbb{P}_s(T_{n;c}) = \mathbb{E}_l \left(\frac{c}{c+1} \right)^l$$

Using this expression for the steady-state success probability, we can show that if Λ is uniform (i.e., $\lambda_{st} = 1/n(n-1)$ for all s, t), the steady-state success probability for tree networks has a lower bound of $\Theta(c/n^2)$ and an upper bound of $\Theta(c/(c+1))$. The lower bound is attained for line networks and the upper bound for star networks.

3.2.2 Cycles

Next we derive the steady-state success probability for cycle graphs. Let $G_{n;c}^\circ$ be a cycle graph of $n > 2$ nodes where each edge has a total credit capacity of c . Let $\mathcal{M}(G_{n;c}^\circ, \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ . Then $\mathcal{M}(G_{n;c}^\circ, \Lambda)$ will have a uniform steady-state distribution over the cycle-reachable equivalence classes of $G_{n;c}^\circ$. We use this fact to show the following result:

Theorem 10. *If $\mathcal{M}(G_{n;c}^\circ, \Lambda)$ is an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in $G_{n;c}^\circ$, then the steady-state transaction success probability, $\mathbb{P}_s(G_{n;c}^\circ)$, is given by*

$$\mathbb{P}_s(G_{n;c}^\circ) = \mathbb{E}_l \frac{r^l + r^{n-l} - 2r^n}{1 - r^n}$$

where $r = c/(c+1)$ and l is the number of edges between a pair of transacting nodes.

The proof involves counting the total number of equivalence classes of $G_{n;c}^\circ$ and the number of classes which allow a transaction between a pair of nodes that have l edges between them. See Appendix A for the proof.

3.2.3 Complete Graphs

Let $K_{n;c}$ be a complete graph over n nodes such that each edge has total capacity c . From Theorem 8, we know that the bankruptcy probability of a node in $K_{n;c}$ is $\Theta(1/nc)$. Next we will show that the steady-state transaction failure probability between two nodes in $K_{n;c}$ is also $\Theta(1/nc)$. Note that a node being bankrupt is a sufficient but not necessary condition for a transaction with that node as the payer to fail.

Theorem 11. *Let $\mathcal{M}(K_{n;c}, \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in $K_{n;c}$. Then the steady-state transaction failure probability in $K_{n;c}$ is $\Theta(1/nc)$.*

See Appendix B for the proof. We will compare this probability with that in an equivalent centralized model to show that liquidity in credit networks with this topology is comparable to the centralized system.

3.2.4 Erdős-Rényi Networks

Here we prove that under a symmetric transaction regime, the steady-state probability that a node in a Erdős-Rényi network, $G_c(n, p)$, will go bankrupt is $\Theta(1/(npc))$. Note that $G_c(n, p)$ is a random graph over n nodes where every edge is present with probability p and has capacity c if present.

Theorem 12. *Consider a $G_c(n, p)$ graph where $p > 8(\ln cn)/n$ (just slightly more than the connectivity threshold of $2 \ln n/n$). Let $\mathcal{M}(G_c(n, p), \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in $G_c(n, p)$. Fix a vertex v . The steady-state probability that v goes bankrupt is $\Theta(1/(npc))$.*

Proof. From Theorem 8, we know that the bankruptcy probability of v is at most $1/(d_v + 1)$. We need to show that in expectation this term is $\Theta(1/(npc))$. By Chernoff bounds, $\mathbb{P}[d_v \leq npc/2] \leq e^{-np/8} \leq e^{-\ln cn} = 1/(cn)$, which bounds $\mathbb{E}_G[1/(d_v + 1)]$ as follows:

$$\begin{aligned} \mathbb{E}_G[1/(d_v + 1)] &\leq \mathbb{P}[d_v < npc/2] \cdot 1 + \mathbb{P}[d_v \geq npc/2] \frac{1}{npc/2} \\ &\leq \frac{1}{cn} + \frac{2}{npc} = \Theta\left(\frac{1}{npc}\right) \end{aligned}$$

□

We conjecture that the steady-state failure probability of transactions between nodes in a $G_c(n, p)$ network under a symmetric transaction regime is also $\Theta(1/(npc))$. This is based on heuristic calculations using the generating functions for complete graphs with pc instead of c , as well as on simulation results (see Section 4.2).

Conjecture 1. *Let $G_c(n, p)$ be a connected credit network (i.e., $p > \ln n/n$) and let $\mathcal{M}(G_c(n, p), \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in $G_c(n, p)$. Then, the steady-state transaction failure probability in $G_c(n, p)$ is $\Theta(1/(npc))$.*

As with complete graphs, we will also compute this probability in an equivalent centralized model and argue that this conjecture, if true, would imply that the failure probability in $G_c(n, p)$ credit networks is within a constant factor of the equivalent centralized model.

3.2.5 Barabási-Albert Preferential Attachment Networks

We state a conjecture regarding the steady-state transaction failure probability in power-law graphs constructed using the Barabási-Albert preferential attachment model. The Barabási-Albert model is a evolving random graph model where each arriving node creates edges to existing nodes with probability proportional to their degrees. It is parameterized by the number of nodes in the network, n , and the number of edges d that each new node creates.

Conjecture 2. *Let $G_c^{BA}(n, d)$ be a random graph over n nodes created using the Barabási-Albert preferential attachment model where each arriving node creates d edges and each edge has total credit capacity c . Let $\mathcal{M}(G_c^{BA}(n, d), \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over nodes in $G_c^{BA}(n, d)$. Then, the steady-state transaction failure probability in $G_c^{BA}(n, d)$ is $\Theta(1/dc)$.*

Note that we conjecture the transaction failure probability to be independent of network size n . This conjecture is supported by simulation results (see Section 4).

3.3 Centralized Currency Infrastructure

Since we introduced credit networks as an alternative to a centralized currency infrastructure, a natural question arises: how does the steady-state failure probability in a credit network compare with that in a centralized infrastructure? Next we address this question.

As described previously, a centralized currency infrastructure can be modeled as a tree with a bank at its root and individuals as leaves. However, this setup is different from a star network in that the root does not participate in any transactions; all transactions occur between leaf nodes, i.e., along paths of length two. Let the centralized infrastructure be represented by a network \mathcal{G}^* with a root node r representing the bank, and n leaf nodes representing individuals. Let $c_{ru} < \infty$ be the credit extended by bank r to node u . Note that since the bank issues a common currency a leaf node trusts the bank for an infinite sum of money, i.e., $c_{ur} = \infty$. A successful transaction of unit value leads to an increase by one of the payee's credit and a decrease by one of the payer's credit; the total credit extended to leaf nodes in the system stays constant. We can prove that if the total credit extended to leaf nodes is C , then the steady-state transaction failure probability in a centralized currency system is $(n - 1)/(C + n - 1)$.

Theorem 13. *Let $\mathcal{M}(\mathcal{G}^*, \Lambda)$ be an ergodic Markov chain induced by a symmetric transaction matrix Λ over a centralized payment network \mathcal{G}^* with n leaf nodes. Let C be the total credit extended to the leaf nodes. Then the steady-state transaction failure probability in \mathcal{M} is $(n - 1)/(C + n - 1)$.*

Proof. Since the number of leaf nodes is n and the total credit extended to leaf nodes is C , the Markov chain has $\binom{C+n-1}{n-1}$ states: this is the number of ways of distributing C objects among n people. In the centralized system, unlike credit networks, a transaction fails if and only if the payer in the transaction, say node u , is bankrupt, *i.e.*, $c_{ru} = 0$. Therefore, the transaction failure probability is equal to the probability that a node is bankrupt. The number of states in which a given node is bankrupt is given by $\binom{C+n-2}{n-2}$: this is the number of ways of distributing the total credit among the remaining $n - 1$ leaf nodes. From Corollary 6, we know that \mathcal{M} has a uniform steady-state distribution over its states. Therefore, the steady-state probability that a node is bankrupt is given by

$$\frac{\binom{C+n-2}{n-2}}{\binom{C+n-1}{n-1}} = \frac{n-1}{C+n-1}$$

□

So if $\bar{c} = C/n$ is the average credit extended to a node in \mathcal{G}^* , then the node bankruptcy probability of a node (as well as the transaction failure probability) is $\Theta(1/\bar{c})$. Note that since the Markov chain has a uniform distribution over states of the network, the steady-state transaction failure probability is independent of the initial state (*i.e.*, the initial credit of each node in the network).

3.3.1 Equivalence between Credit Networks and Centralized Currency Infrastructure

In order to be able to compare liquidity in the two models, we need to define the notion of equivalence between the two models. Given a credit network $G = (V, E)$ with n nodes, we construct an equivalent centralized currency infrastructure $\mathcal{G}(G)$ as follows: \mathcal{G} has nodes $V \cup \{r\}$ (where r is the root node) and a bidirectional edge between each node u and the root r . For each node u in G , the total credit extended to it by the bank in the centralized model, c_{ru} , is given by $c_{ru} = \sum_{v \in V} c_{vu}$ where c_{vu} is the initial credit extended by v to u before the nodes start transacting. And finally, for all nodes u , we set $c_{ur} = \infty$. This ensures that each node has access to the same amount of credit in both models.

3.3.2 Liquidity Comparison

The procedure to convert a credit network in to an equivalent centralized currency system allows us to compare liquidity in credit networks with that in an equivalent currency system. For the credit network topologies that we analyzed in Section 3.2, we state the steady-state failure probability for their equivalent centralized systems under a symmetric, ergodic transaction regime by invoking Theorem 13.

Cycle Graph For a cycle graph with n nodes where each edge has capacity c , the total capacity is nc . So, the steady-state failure probability in the equivalent centralized system is $(n-1)/(nc+n-1) = \Theta(1/c)$. On the other hand, in a uniform transaction regime, credit networks on cycle graphs have a steady-state failure probability of $1 - \Theta(c/n)$ (see Appendix A). However, for graphs with high expansion, the failure probability in credit networks is within a constant factor of that in equivalent centralized systems (see below).

Complete Graph For a complete graph on n nodes where the total credit capacity between any pair of nodes is c , the total credit capacity in the system is Nc , where $N = \binom{n}{2}$. So, the steady-state failure probability in the equivalent centralized system is $(n-1)/(Nc+n-1) = \Theta(1/n)$, the same as that in the original credit network (Theorem 11).

Erdős-Rényi Graph For an Erdős-Rényi network, $G_c(n, p)$, the expected total credit capacity in the network is Npc , where $N = \binom{n}{2}$. Therefore, the steady-state failure probability in the equivalent centralized system is $(n-1)/(Npc+n-1) = \Theta(1/n)$. Thus, if Conjecture 1 is true, the failure probability in $G_c(n, p)$ credit networks would be within a constant factor of that in the equivalent centralized system.

Barabási-Albert Graph For a Barabási-Albert preferential attachment network, $G_c^{BA}(n, d)$, the total credit capacity in the network is ndc (*i.e.*, nd edges each having total credit c). Therefore, the steady-state failure probability in the equivalent centralized system is $(n-1)/(ndc+n-1) = \Theta(1/d)$. Like Erdős-Rényi networks, if the conjecture about steady-state failure probability for $G_c^{BA}(n, d)$ credit networks (Conjecture 2) is true, then, credit networks with this topology too will have a steady-state failure probability within a constant factor of that in an equivalent centralized system.

4 Simulations

Next we present the results of simulating repeated transactions on credit networks from two well-studied families of random graphs: $G_c(n, p)$ (this is a Erdős-Rényi random graph with edge orientation chosen randomly and each edge assigned capacity c) and power-law graphs generated using the Barabási-Albert preferential attachment (BA) model [BA99]. The Barabási-Albert model is an evolving random graph model where each new node creates d edges to existing nodes with probability proportional to their degrees. After constructing the graph, we assigned edge orientations randomly and set the capacity of each edge to c . The goal of the simulations was to understand how the steady-state success probability in these networks depends on key network parameters and to determine what values of network parameters yield a “high enough” success probability. In this sense, the qualitative results obtained from running simulations complement the asymptotic analysis described in the previous section. For $G_c(n, p)$ the parameters of interest are n, p and c . However, since varying either n or p changes the density of the graph (*i.e.*, average node degree), we also ran simulations where we changed both n and p such that np was held constant. For BA graphs, the parameters of interest are the number of nodes (n), the number of edges each node creates (d), and credit capacity (c) of each edge.

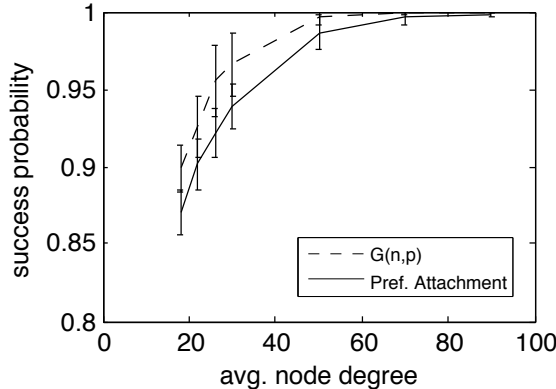


Figure 2: Effect of varying graph density (average node degree) on the steady-state success probability in $G_c(n, p)$ and $G_c^{BA}(n, d)$ graphs with $n = 200$ and $c = 1$.

A simulation run consisted of constructing a network with a given set of parameters and doing repeated transactions on it. At each time step, we chose a node pair (s, t) with uniform probability and tried to route a unit payment from s to t via the shortest feasible path. If there was a path from s to t , we routed the payment and modified edge capacities along the path. Otherwise, we counted the transaction as a failure. We repeated this until the success-rate (*i.e.* fraction of transactions that succeed) in two consecutive time windows was within ε of each other. We used a window size of 1000 time steps and $\varepsilon = 0.002$. When the process converged, we measured the success rate at convergence, *i.e.*, the total number of successful transactions divided by the total number of time steps. We called this the steady-state success probability for that set of network parameters. For each combination of network parameter values, we ran simulations on 100 graphs constructed from the same distribution and computed the average and the standard deviation in success-rate over the 100 runs. In order to understand how the success probability depends upon various network parameters, we repeated the set of 100 runs for different values of one of the parameters, keeping the rest constant.

4.1 Effect of Variation in Network Density

In order to study the effect of density, we fixed $n = 200$ and $c = 1$ for both graph families. For $G_c(n, p)$ we varied p from 0.18 to 0.45 and for BA graphs we varied d from 18 to 45. The average node degree of PA graphs is equal to $2d$ whereas for $G_c(n, p)$ the expected node degree is given by $(n - 1)p$. We picked d corresponding to each p such that $2d \approx (n - 1)p$ so that we can compare the plots for the two graph families.

Figure 2 shows the success probability as a function of average node degree for $G_c(n, p)$ and BA graphs when np is sufficiently larger than $\ln n$ so that the $G_c(n, p)$ graph is connected. The success rate for both graph families is concave, non-decreasing in p . Note that for networks of 200 nodes with edge capacity of 1 and an average node degree as low as 25, the success probability for both topologies is in excess of 0.9. Moreover, the success probability quickly reaches a point of diminishing returns where making the network denser does not help significantly. This is useful since it means that, in practice, the network does not need to be very dense in order to route payments with a sufficiently high probability.

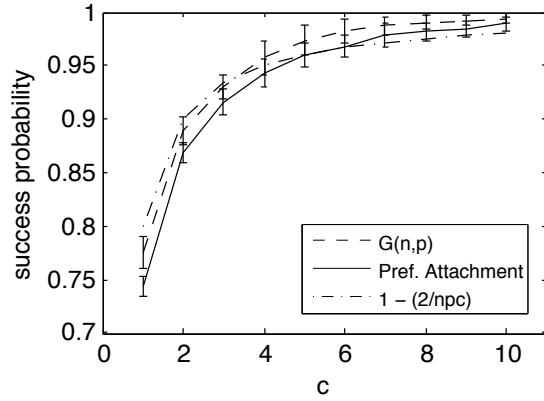


Figure 3: Effect of varying credit capacity on the steady-state success probability in $G_c(n, p)$ and $G_c^{BA}(n, d)$ graphs with $n = 100$ and avg. node degree = 10.

4.2 Effect of Variation in Credit Capacity

In order to study the effect of varying credit capacity in the network, we fixed $n = 100$ for both graphs. We set $p = 0.10$ for $G_c(n, p)$ and $d = 5$ for BA graphs, thereby ensuring that their densities were equal. This choice of p and d also ensured that all the graphs we constructed were connected. We varied c from 1 to 10. Figure 3 shows the mean and standard deviation in the steady-state success-rate for both graph families as we vary c . In addition to the two plots, we also plot $1 - (2/npc)$ as a function of c . That plot tracks those for $G_c(n, p)$ as well as for BA fairly closely and is further evidence in support of our conjecture that the steady-state failure probability in $G_c(n, p)$ graphs is $\Theta(1/npc)$ and in $G_c^{BA}(n, d)$ is $\Theta(1/dc)$ and consequently, credit networks with these topologies are within a constant-factor of the equivalent centralized currency system in terms of liquidity.

4.3 Effect of Variation in Network Size

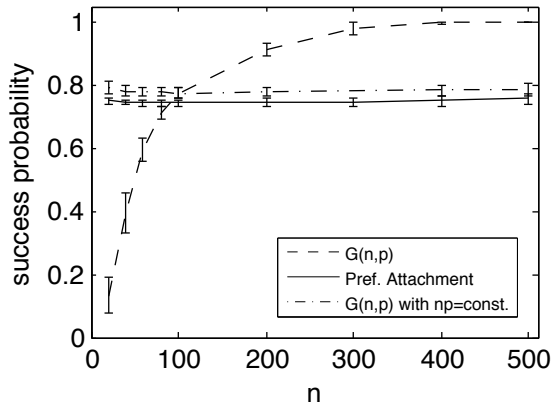


Figure 4: Effect of varying network size (number of nodes) on the steady-state success probability in $G_c(n, p)$ and $G_c^{BA}(n, d)$ graphs with $c = 1$ and avg. node degree = 10.

In order to study the effect of varying network size for PA graphs, we fixed $d = 5$ and $c = 1$ and varied n from 20 to 500. However, the same approach does not work for $G_c(n, p)$ graphs since if we fix p and c and varied n , it would also vary the average node degree. So, we ran two sets of simulations for $G_c(n, p)$ graphs: one where we varied n keeping $p = 0.10$ and $c = 1$ and the other where we fixed $c = 1$ and varied both n and p such that np was kept constant at 10.

Figure 4 shows the effect of varying network size on the success probability of both graph families. The noteworthy observation is that if the density of the network is kept constant, then network size has no effect on the success rate for either graph family; the success rate remains nearly constant at about 0.75 for BA graphs and about 0.78 for $G_c(n, p)$ graphs. The same pattern is also observed if np is held constant at 20 instead of 10 (so it is not an artifact of that specific value of np). This is another observation in support of our conjecture about the steady-state failure probability in $G_c(n, p)$ and BA graphs. If,

however, np is not held constant, then success probability is again a concave non-decreasing function of n (much the same way as it is of p).

5 Conclusion & Open Problems

In this paper we generalized the credit network model by allowing every node to print its own currency, and formulated and studied the question of long-term liquidity in this network under a simple model of repeated transactions. Using the notion of cycle-reachability, we showed that routing payments in credit networks has a *path-independence* property. We also showed that the Markov chain induced by a symmetric transaction regime has a uniform steady-state distribution over the cycle-reachable equivalence classes. We used this fact to derive the node bankruptcy probability in general graphs and the transaction failure probability in a number of network topologies. We showed that, except in cycles, these probabilities are within a constant factor of the corresponding values for an equivalent centralized payment infrastructure. Our analysis and simulations show that for a number of well-known graph families, the steady-state failure probability under reasonable transaction regimes is comparable to that in equivalent centralized currency systems. Thus, in return for the robustness and decentralized properties, this model does not lose much liquidity compared to a centralized model.

5.1 Open Problems related to Liquidity

Our results on liquidity can be extended in a number of ways. One question that arises from our analysis is whether credit network topologies with a high expansion have comparable liquidity with equivalent centralized systems. In general, can the steady state success probability in general networks be expressed in terms of some key network properties such as density or expansion? Another direction is to understand how success-rate and rate of convergence vary under various transaction regimes. A natural regime to study would be where the probability that a node is a payer (payee) is proportional to the node's in-degree (out-degree). We also leave unanswered the question of mixing times of the Markov chain: our simulations on $G_c(n, p)$ and BA graphs converged relatively quickly but we do not have analytical results on the mixing time of the Markov chain and how it depends on credit capacity and network topology. One can also ask how fault-tolerant credit networks are under various models of node failures. For example, if we remove k nodes (and all their edges) from the network, how would liquidity change in the worst-case, and what network topologies are most robust against such failures? If node failures cascade through the network, how quickly would the network lose liquidity? Another question worth investigating is the effect of transaction routing costs on liquidity. It appears that a mathematical analysis of a transaction routing model where intermediate nodes charged a routing fee would require an entirely new approach since it would invalidate the cycle-reachability relation that forms the basis of our results.

5.2 Further Understanding the Model

This model presents a number of interesting directions for further inquiry. The first is the question of pricing of currencies. Let r_{uv} be the exchange ratio for converting u 's currency into v 's (*i.e.*, 1 unit of u 's currency is r_{uv} units of v 's currency). The following two properties of these conversion ratios are necessary for the generalized credit network model to be well-defined:

- *Conservation:* For any $S \subseteq V$, where $S = \{u_1, u_2, \dots, u_k\}$, $\prod r_{u_1 u_2} r_{u_2 u_3} \dots r_{u_{k-1} u_k} r_{u_k u_1} = 1$. In words, this property means that currency is conserved along any cycle of payments in the network.
- *Common Knowledge:* The conversion ratios are common knowledge for all nodes in V . This means that nodes along a path through which payment is being routed can compute the outflow given the inflow.

However, these properties do not guarantee that it is incentive-compatible for nodes to route payments. For example, it is natural to believe that the currency of a node that lies on the path between a number of (source, destination) pairs should be more expensive compared to one that has, say, only one neighbor. More generally, can we endow the nodes in the network with a model of rationality that will govern how much credit nodes extend each other, how that changes over time, and how nodes negotiate currency exchange ratios as a function of transaction rates, success probabilities, network topology, etc.?

The second broad question is how to apply this model to various practical applications, some of which we mention in the introduction. Of particular interest are analytical results as well as empirical or simulation-based studies that evaluate the effectiveness of this model in settings such as IP routing, peer-to-peer systems, viral marketing, etc. compared to existing models.

Another potential direction is to consider credit networks as enabling a trust-based market where nodes buy and sell goods using trusted currency instead of a common currency. Then we can ask the same questions that have been asked of various traditional market models: Does the market clear? At what prices? Can we efficiently find them? How do equilibrium prices in this model compare with those in models that assume a common currency? Answers to these questions will provide insights into

whether and when can we trade off the robustness properties and the decentralized nature of this model with the well-understood characteristics of traditional markets.

6 Acknowledgements

This research was supported in part by NSF award IIS-0904325. Part of the research was also sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- [AJ06] Christina Aperjis and Ramesh Johari. A Peer-to-Peer System as an Exchange Economy. In *Proceedings of the 2006 Workshop on Game theory for Communications and Networks*, GameNets '06, New York, NY, USA, 2006. ACM. Available from: <http://doi.acm.org/10.1145/1190195.1190206>, doi:<http://doi.acm.org/10.1145/1190195.1190206>.
- [BA99] A. L. Barabasi and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999. doi:[10.1126/science.286.5439.509](https://doi.org/10.1126/science.286.5439.509).
- [DB05] Dimitri B. DeFigueiredo and Earl T. Barr. Trustdavis: A non-exploitable online reputation system. In *CEC '05: Proceedings of the 7th IEEE International Conference on E-Commerce Technology*, pages 274–283, Washington, DC, USA, 2005. doi:[10.1109/ICECT.2005.98](https://doi.org/10.1109/ICECT.2005.98).
- [GGMP04] Zoltán Gyöngyi, Hector Garcia-Molina, and Jan Pedersen. Combating web spam with trustrank. In *VLDB '04: Proceedings of the 30th International Conference on Very Large Databases*, pages 576–587. ACM, 2004. Available from: <http://portal.acm.org/citation.cfm?doid=1316689.1316740>.
- [Gio07] Emeric Gioan. Enumerating Degree Sequences in Digraphs and a Cycle-cocycle Reversing System. *European Journal of Combinatorics*, 28(4):1351–1366, 2007. doi:[10.1016/j.ejc.2005.11.006](https://doi.org/10.1016/j.ejc.2005.11.006).
- [GKRT04] R. Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of Trust and Distrust. In *WWW '04: Proceedings of the 13th International Conference on World Wide Web*, pages 403–412, New York, NY, USA, 2004. ACM. doi:[10.1145/988672.988727](https://doi.org/10.1145/988672.988727).
- [GMR⁺07] Arpita Ghosh, Mohammad Mahdian, Daniel M. Reeves, David M. Pennock, and Ryan Fugger. Mechanism Design on Trust Networks. In *WINE '07: Proceedings of the 3rd international workshop on Internet and Network Economics*, pages 257–268, 2007. doi:[10.1007/978-3-540-77105-0_25](https://doi.org/10.1007/978-3-540-77105-0_25).
- [JRBT11] Matthew O. Jackson, Tomás Rodríguez Barraquer, and Xu Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, pages 1–52, 2011. Available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1657130.
- [KFH07] Ian A. Kash, Eric J. Friedman, and Joseph Y. Halpern. Optimizing Scrip Systems: Efficiency, Crashes, Hoarders, and Altruists. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 305–315, 2007.
- [KMRS09] Dean Karlan, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. Trust and Social Collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, August 2009. Available from: <http://ideas.repec.org/a/tpr/qjecon/v124y2009i3p1307-1361.html>.
- [KW81] Daniel J. Kleitman and Kenneth J. Winston. Forests and Score Vectors. *Combinatorica*, 1(1):49–54, 1981. doi:[10.1007/BF02579176](https://doi.org/10.1007/BF02579176).
- [LHK10] Jure Leskovec, Daniel P. Huttenlocher, and Jon M. Kleinberg. Signed Networks in Social Media. In *CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pages 1361–1370. ACM, 2010. doi:[10.1145/1753326.1753532](https://doi.org/10.1145/1753326.1753532).

- [LHL⁺10] Zhengye Liu, Hao Hu, Yong Liu, Keith W. Ross, Yao Wang, and Markus Mobius. P2P Trading in Social Networks: The Value of Staying Connected. In *IEEE INFOCOMM '10: Proceedings of the 29th IEEE International Conference on Computer Communications*, pages 1–9, March 2010. doi:10.1109/INFCOM.2010.5462077.
- [MPGD08] Alan Mislove, Ansley Post, Krishna P. Gummadi, and Peter Druschel. Ostra: Leverging Trust to Thwart Unwanted Communication. In *Proceedings of the 5th Symposium on Networked Systems Design and Implementation (NSDI'08)*, April 2008. Available from: http://www.usenix.org/events/nsdi08/tech/full_papers/mislove/mislove.pdf.
- [PBMW99] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank Citation Ranking: Bringing Order to the Web. Technical report, Stanford InfoLab, 1999. Available from: <http://ilpubs.stanford.edu:8090/422/>.
- [RS09] Paul Resnick and Rahul Sami. Sybilproof Transitive Trust Protocols. In *EC '09: Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 345–354, 2009. doi:10.1145/1566374.1566423.
- [RSK07] Daniel M. Reeves, Bethany M. Soule, and Tejaswi Kasturi. Yootopia! Technical report, Yahoo! Research, 2007. Available from: <http://www.yootles.com/yootles.pdf>.
- [VCS03] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing. In *Workshop on the Economics of Peer-to-Peer Systems*, 2003. Available from: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.5.6462>.

A Success Probability in Cycles

Proof of Theorem 10. Since \mathcal{M} has a uniform steady-state distribution over the cycle-reachable equivalence classes, the steady-state success probability between a pair of nodes separated by l edges in $G_{m,c}^o$ is simply the fraction equivalence classes which allow a transaction between that pair of nodes.

Note that each equivalence class can be characterized by a state in which at least one edge in $G_{c,m}^o$ has zero capacity in the counterclockwise direction (any state where all edges are bidirectional is cycle-reachable from one where at least one edge is unidirectional). Assume that the nodes are numbered $1, \dots, m$ and the edge between nodes j and $j + 1$ is numbered j . Let the j th edge be the first edge that has zero capacity in the counterclockwise direction. Then the first $j - 1$ edges can each be in one of c states, and the remaining $m - j$ edges can each be in one of $c + 1$ states. Thus, the total number of states where one or more edges have zero capacity in the counterclockwise direction is given by:

$$\begin{aligned} \sum_{j=1}^m c^{j-1}(c+1)^{m-j} &= (c+1)^{m-1} \sum_{j=1}^m \left(\frac{c}{c+1}\right)^{j-1} \\ &= (c+1)^{m-1} \frac{1-r^m}{1-r} \end{aligned}$$

where $r = c/(c+1)$. Now fix nodes s and t that are a distance l apart in the clockwise direction from s to t . Lets call the path of length l between s and t , \mathcal{P}_1 , and that of length $m - l$, \mathcal{P}_2 . A successful transaction between s and t along \mathcal{P}_1 requires edges in \mathcal{P}_1 to have capacity in the counterclockwise direction, whereas that along \mathcal{P}_2 requires edges in \mathcal{P}_2 to have capacity in the clockwise direction. If the j th edge lies in \mathcal{P}_1 then there is no path of length l between s and t . On the other hand, a path of length $m - l$ between s and t requires each of the $m - l$ edges to be in one of c states. Thus, the number of equivalence classes that only have a path of length $m - l$ between s and t is

$$\begin{aligned} \sum_{j=1}^l c^{j-1}(c+1)^{l-j}c^{m-l} &= (c+1)^{l-1}c^{m-l} \sum_{j=1}^l \left(\frac{c}{c+1}\right)^{j-1} \\ &= c^{m-l}(c+1)^{l-1} \frac{1-r^l}{1-r} \end{aligned}$$

where $r = c/(c+1)$. However, if the j th edge is in \mathcal{P}_2 , it implies that all edges in \mathcal{P}_1 have positive capacity in the counterclockwise

direction. Thus, the number of equivalence classes that have a path of length l between s and t is

$$\begin{aligned} \sum_{j=l+1}^m c^{j-1}(c+1)^{m-j} &= c^l(c+1)^{m-l-1} \sum_{j=1}^{m-l} \left(\frac{c}{c+1}\right)^{j-1} \\ &= c^l(c+1)^{m-l-1} \frac{1-r^{m-l}}{1-r} \end{aligned}$$

Note that the above equation includes equivalence classes that have capacity along both \mathcal{P}_1 and \mathcal{P}_2 as well as those that have capacity along only \mathcal{P}_1 .

Thus, the probability that a transaction between nodes separated by l edges will succeed is given by

$$\begin{aligned} &\frac{c^l(c+1)^{m-l-1}(1-r^{m-l}) + c^{m-l}(c+1)^{l-1}(1-r^l)}{(c+1)^{m-1}(1-r^m)} \\ &= \frac{r^l(1-r^{m-l}) + r^{m-l}(1-r^l)}{1-r^m} = \frac{r^l + r^{m-l} - 2r^m}{1-r^m} \end{aligned}$$

Taking expectation over the length l of the transaction path gives the result. \square

As a corollary, we can show that the steady-state success probability in cycle graphs under a uniform transaction regime is $\Theta(c/n)$.

Corollary 14. *Let $G_{n;c}^\circ$ be a cycle graph of $n > 2$ nodes where each edge has a total credit capacity of c . If $\mathcal{M}(G_{n;c}^\circ, \Lambda)$ is an ergodic Markov chain induced by a uniform transaction matrix Λ over nodes in $G_{n;c}^\circ$, then the steady-state transaction success probability, $\mathbb{P}_s(G_{n;c}^\circ)$, is $\Theta(c/n)$.*

Proof. $\mathbb{P}_s(G_{n;c}^\circ)$ under a symmetric transaction regime is given by

$$\mathbb{P}_s(G^\circ) = \mathbb{E}_l \frac{r^l + r^{n-l} - 2r^n}{1-r^n} \quad (\text{where } r = c/(c+1))$$

Therefore, under a uniform transaction regime, as $n \rightarrow \infty$

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}_s(G^\circ) &\leq \lim_{n \rightarrow \infty} \sum_{l=1}^{n/2} \frac{2n}{n(n-1)} \frac{r^l + r^{n-l} - 2r^n}{1-r^n} \\ &\leq \lim_{n \rightarrow \infty} \frac{2}{(n-1)(1-r^n)} \sum_{l=1}^{n/2} (r^l + r^{n-l}) \\ &= \lim_{n \rightarrow \infty} \frac{2}{(n-1)(1-r^n)} \left(\frac{r(1-r^{n/2}) + r^{n/2}(1-r^{n/2})}{1-r} \right) \\ &= \frac{2r}{(n-1)(1-r)} = \Theta(c/n) \end{aligned}$$

\square

B Failure Probability in Complete Graphs

Proof of Theorem 11. We know that there is a bijection between equivalence classes of a graph and labeled forests [Gio07, KW81], and the number of forests, $a_{n,c}$, in $K_{n;c}$ converges to [KW81]

$$a_{n,c} \sim e^{1/2c} c^{n-1} n^{n-2} \quad (1)$$

where $f_n \sim g_n$ means $\lim_{n \rightarrow \infty} f_n/g_n = 1$. By Cayley's formula there are n^{n-2} trees on n labeled nodes, so $a_{n,1} \geq n^{n-2}$. Any forest has at most $n-1$ edges, so with c labeled copies of each edge $c^{n-1} n^{n-2} \leq a_{n,c} \leq c^{n-1} a_{n,1}$. Equation (1) implies that there exists a constant n_0 such that for all $n \geq n_0$, $a_{n,1} \leq 2e^{1/2} n^{n-2}$. Combining these equations, for all $n \geq n_0$,

$$c^{n-1} n^{n-2} \leq a_{n,c} \leq 2e^{1/2} c^{n-1} n^{n-2}$$

Fix two nodes u and v in $K_{n;c}$, and consider a cut between u and v such that the partition containing u has k nodes and that containing v has $n - k$ nodes. Assume that no payment can be made from u 's partition to v 's partition. Then, the number of equivalence classes containing this cut is at most $a_{k,c}a_{n-k,c}$. Summing over all cuts separating u and v in $K_{n;c}$, we upper bound the transaction failure probability between nodes u and v as

$$\frac{1}{a_{n,c}} \sum_{k=1}^{n-1} \binom{n-2}{k-1} a_{k,c} a_{n-k,c} = 2 \sum_{k=1}^{n_0-1} \binom{n-2}{k-1} \frac{a_{k,c} a_{n-k,c}}{a_{n,c}} + \sum_{k=n_0}^{n-n_0} \binom{n-2}{k-1} \frac{a_{k,c} a_{n-k,c}}{a_{n,c}}$$

Bounding the first term,

$$\begin{aligned} \sum_{k=1}^{n_0-1} \binom{n-2}{k-1} \frac{a_{k,c} a_{n-k,c}}{a_{n,c}} &= O\left(\sum_{k=1}^{n_0-1} \frac{n^{k-1}}{(k-1)!} \frac{(c^{k-1} a_{k,1})(c^{n-k-1} (n-k)^{n-k-2})}{c^{n-1} n^{n-2}}\right) \\ &= O\left(\frac{1}{nc} \sum_{k=1}^{n_0-1} \frac{a_{k,1}}{(k-1)!} \left(\frac{n-k}{n}\right)^{n-k-2}\right) = O\left(\frac{1}{nc}\right) \end{aligned}$$

using that $a_{k,1} = O(1)$ when $k \leq n_0$. For the second term,

$$\begin{aligned} &\sum_{k=n_0}^{n-n_0} \binom{n-2}{k-1} \frac{a_{k,c} a_{n-k,c}}{a_{n,c}} \\ &= O\left(\sum_{k=n_0}^{n-n_0} \binom{n-2}{k-1} \frac{(c^{k-1} k^{k-2})(c^{n-k-1} (n-k)^{n-k-2})}{c^{n-1} n^{n-2}}\right) \\ &= O\left(\frac{1}{nc} \sum_{k=n_0}^{n-n_0} \binom{n-2}{k-1} \frac{k^{k-2} (n-k)^{n-k-2}}{n^{n-3}}\right) \\ &= O\left(\frac{1}{nc} \sum_{k=1}^{n-1} \frac{n!}{n(n-1)} \frac{k(n-k)}{k!(n-k)!} \frac{k^k (n-k)^{n-k}}{n^n} \frac{n^3}{k^2 (n-k)^2}\right) \\ &= O\left(\frac{1}{nc} \sum_{k=1}^{n-1} \frac{n}{(n-k)k} \frac{n!}{n^n} \frac{k^k (n-k)^{n-k}}{k! (n-k)!}\right) \end{aligned} \tag{2}$$

Instead of using the standard Stirling's approximation, we will use the more precise *Stirling Series*:

$$n! = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n \left(1 + \frac{1}{12n} + \frac{1}{288n^2} - \frac{139}{51840n^3} \pm \dots\right)$$

Therefore, $n!$ can be written as

$$n! = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n c_n, \text{ where } 1 \leq c_n \leq 2 \text{ when } n \geq 1$$

Substituting this into (2) and simplifying, we get

$$\frac{1}{nc} \sum_{k=1}^{n-1} \frac{n}{(n-k)k} \frac{n!}{n^n} \frac{k^k (n-k)^{n-k}}{k! (n-k)!} = O\left(\frac{1}{nc} \sum_{k=1}^{n-1} \left(\frac{n}{(n-k)k}\right)^{3/2}\right)$$

Either k or $n - k$ is at least $n/2$, so

$$\sum_{k=1}^{n-1} \left(\frac{n}{(n-k)k}\right)^{3/2} \leq 2 \sum_{k=1}^{(n-1)/2} \left(\frac{2}{k}\right)^{3/2} = O(1)$$

Summarizing,

$$\frac{1}{a_{n,c}} \sum_{k=1}^{n-1} \binom{n-2}{k-1} a_{k,c} a_{n-k,c} = O\left(\frac{1}{nc}\right)$$

Also, since the node bankruptcy probability in $K_{n;c}$ is $\Theta(1/nc)$, the transaction failure probability is $\Omega(1/nc)$. This completes the proof. \square