

On the inapproximability of minimizing cascading failures under the deterministic threshold model



Lei Nie^{a,b,*}, Jingjie Liu^{a,b}, Haicang Zhang^{a,b}, Zhiwei Xu^a

^a Institute of Computing Technology, Chinese Academy of Sciences, China

^b University of Chinese Academy of Sciences, China

ARTICLE INFO

Article history:

Received 8 July 2012

Received in revised form 7 October 2013

Accepted 21 October 2013

Available online 25 October 2013

Communicated by J. Torán

Keywords:

Combinatorial problems

Inapproximability

Cascading failures

Deterministic threshold model

ABSTRACT

Given a directed graph G and a threshold $L(r)$ for each node r , the rule of deterministic threshold cascading is that a node r fails if and only if it has at least $L(r)$ failed in-neighbors. The cascading failure minimization problem is to find at most k edges to delete, such that the number of failed nodes is minimized. We prove an $n^{1-\epsilon}$ inapproximability result for the general case and a $\frac{1}{2}n^\epsilon$ inapproximability result for the special case with the maximum threshold of 1.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Cascading failures are ubiquitous in various types of networks. For example, bankruptcies of several banks can shock other banks through a financial network and lead to a global financial crisis [1]. The spread of cascading failures is highly relevant to the topology of the propagation network [2]. There are several methods, such as blocking edges and vaccinating nodes, which can modify networks more resilient to cascading failures. This paper focuses on minimizing the number of cascading failures by deleting a bounded number of edges.

In this paper, we use the concept of *deterministic threshold* to model the cascading processes in a network [2,3]. Let $G = \langle V, E, L \rangle$ be a directed graph with thresholds, where V is a set of nodes, $E \subseteq V \times V$ is a set of directed edges and $L : V \rightarrow \mathbf{N}$ is a threshold function mapping each node to a nonnegative integer. We define the *maximum threshold* $L_{\max} = \max_{r \in V} L(r)$. The rule of deterministic threshold cascading is that a node r fails if and

only if it has at least $L(r)$ failed in-neighbors. At time 0, a cascading process is triggered by nodes with thresholds of 0 (initial failed nodes). At time t , every node r checks whether it has at least $L(r)$ failed in-neighbors. If so, the node r will fail. The state of a node can only switch from alive to failed, but cannot switch from failed to alive. The process stops at time t if and only if there are no new failed nodes. We use $\Phi(V, E, L)$ to denote the set of failed nodes at the end of a cascading process. It can be proved that Φ is uniquely determined. We define the *cascading failure minimization (CFM)* problem as follows.

Problem 1 (Cascading failure minimization). Given a nonnegative integer k (budget) and a directed graph with thresholds $G = \langle V, E, L \rangle$, find at most k edges to delete, such that the number of failed nodes is minimized. That is,

$$\begin{aligned} \min_{E' \subseteq E} & |\Phi(V, E - E', L)| \\ \text{s.t.} & |E'| \leq k. \end{aligned}$$

We study the hardness of the cascading failure minimization problem and prove the first inapproximability

* Corresponding author.

E-mail addresses: nielei@ict.ac.cn (L. Nie), liujingjie@ict.ac.cn (J. Liu), zhanghaicang@ict.ac.cn (H. Zhang), zxu@ict.ac.cn (Z. Xu).

result by a polynomial time reduction from the vertex cover problem.

Theorem 1. *It is NP-hard to approximate the cascading failure minimization problem within a factor of $n^{1-\epsilon}$, for any $0 < \epsilon \leq 1$.*

The maximum threshold is equal to 2 in the above reduction, which motivates us to study the special case with the maximum threshold of 1. We prove the second inapproximability result by an approximation-preserving reduction from the *minimum coverage problem* [4].

Theorem 2. *Unless there exists a probabilistic algorithm for SAT that runs in time 2^{n^δ} on an instance of size n , there is no $\frac{1}{2}n^\epsilon$ -approximation algorithm for the cascading failure minimization problem even if the maximum threshold equals 1, where $\delta > 0$ is an arbitrarily small constant and $\epsilon = 1/2^{O(1/\delta \log(1/\delta))}$.*

The cascading failure minimization problem with the maximum threshold of 1 is a special case of the *quarantining problem* [5] and the *minimum-size bounded-capacity cut problem* [6]. We can get a $\frac{1}{2}n^\epsilon$ inapproximability result of the above two problems as a corollary of **Theorem 2**. In addition, there are several related problems in which the initial failed nodes are not known in advance [7,8]. In contrast to minimize the spread of cascades, the influence maximization problem, motivated by viral marketing, aims to choose a set of k individuals to adopt a new product, such that the number of further adoptions through the cascade in social networks is maximized [9]. Several hardness results and approximation algorithms for this problem were reviewed by Kempe [10].

2. Proof of Theorem 1

We prove **Theorem 1** by a polynomial time reduction from the NP-complete vertex cover problem. Given a vertex cover instance with a vertices, b edges and a non-negative integer k , we can construct a CFM instance as follows. For each edge in the vertex cover instance, we add a node with a threshold of 2 in the CFM instance. We call them edge-nodes. For each vertex in the vertex cover instance, we add a node with a threshold of 1 in the CFM instance. We call them vertex-nodes. If a vertex and an edge are adjacent in the vertex cover instance, we produce a directed edge from the corresponding vertex-node to the corresponding edge-node in the CFM instance. Next, we add a node with a threshold of 0 in the CFM instance, which is called zero-node. We produce directed edges from the zero-node to each vertex-node. We add M nodes with thresholds of 1 in the CFM instance, which are called slave-nodes, where

$$M = (2(a+1)(1+2a+3b))^{\frac{1}{\epsilon}}. \quad (1)$$

We produce directed edges from each edge-node to each slave-node. Finally, let the budget of the CFM instance be k given in the vertex cover instance. It is clear that this reduction can be completed in polynomial time. The length

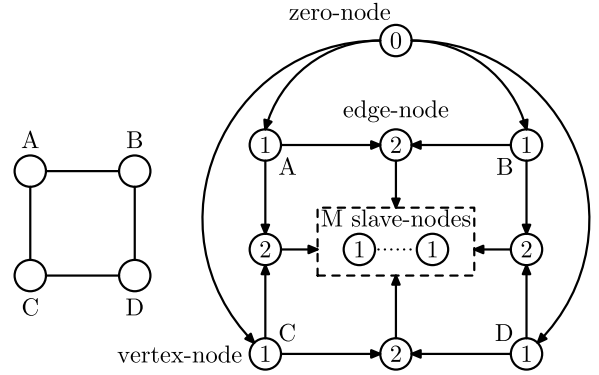


Fig. 1. An example of the reduction from a vertex cover instance (left) to a CFM instance (right).

of the reduced CFM instance is $1 + 2a + 3b + (b + 1)M$, which is the total number of the nodes and edges. **Fig. 1** shows an example of this reduction.

Lemma 1. *Given a CFM instance reduced from a vertex cover instance, any optimal set of edges to delete only contains the edges from the zero-node to vertex-nodes.*

Proof. Because k is less than or equal to the number of vertices, the budget in the CFM instance is less than or equal to the number of edges from the zero-node to vertex-nodes. If there is an optimal set containing an edge from a vertex-node to an edge-node or an edge from an edge-node to a slave-node, we can replace it with an edge from the zero-node to a vertex-node, which will prevent at least one more node from failing. \square

Lemma 2. *If the answer of a vertex cover instance is YES, then the OPT (optimal value) of the reduced CFM instance is $a - k + 1$.*

Proof. Suppose that there is a k vertex cover. We delete the k edges from the zero-node to the vertex-nodes, corresponding vertices of which are in the vertex cover. Because every edge in the vertex cover has one or two endpoints in the vertex cover, all of the edge-nodes and the slave-nodes are alive. The number of failed nodes is $a - k + 1$. According to **Lemma 1**, any optimal set of edges to delete only contains the edges from the zero-node to vertex-nodes. It can save k vertex-nodes at most by deleting k edges from the zero-node to vertex-nodes. Thus, $a - k + 1$ is the OPT. \square

Lemma 3. *If the answer of a vertex cover instance is NO, then the OPT of the reduced CFM instance is greater than M .*

Proof. According to **Lemma 1**, any optimal set of edges to delete only contains the edges from the zero-node to vertex-nodes. Moreover, we cannot find a set of k edges from the zero-node to vertex-nodes in the reduced CFM instance to delete to make sure all edge-nodes are alive, because there is not a k vertex cover. The OPT is therefore greater than M , since all of the M slave-nodes will fail. \square

Lemma 4. *The answer of a vertex cover instance is YES if and only if the $n^{1-\epsilon}$ -approximation value of the reduced CFM instance is less than or equal to M .*

Proof. We use VAL to denote the $n^{1-\epsilon}$ -approximation value of the reduced CFM instance. Suppose that $VAL \leq M$. Thus, $OPT \leq VAL \leq M$. We can conclude that there exists a k vertex cover; otherwise the OPT must be greater than M according to Lemma 3.

Suppose that there is a k vertex cover. According to Lemma 2, $OPT = a - k + 1$. The length of the reduced CFM instance is $1 + 2a + 3b + (b + 1)M$ according to the reduction. Thus,

$$VAL \leq OPT(1 + 2a + 3b + (b + 1)M)^{1-\epsilon}. \quad (2)$$

Deriving from Eq. (1),

$$(1 + 2a + 3b + (b + 1)M)^\epsilon > 2OPT(1 + 2a + 3b). \quad (3)$$

According Eq. (2) and Eq. (3), we can get

$$\begin{aligned} VAL &\leq \frac{OPT(1 + 2a + 3b + (b + 1)M)}{(1 + 2a + 3b + (b + 1)M)^\epsilon} \\ &< \frac{OPT(1 + 2a + 3b + (b + 1)M)}{2OPT(1 + 2a + 3b)} \\ &\leq \frac{1}{2}(1 + M) \\ &\leq M. \quad \square \end{aligned}$$

According to Lemma 4, we have constructed a polynomial time algorithm for the vertex cover problem using an $n^{1-\epsilon}$ -approximation algorithm for the CFM problem. Hence, Theorem 1 is proved.

3. Proof of Theorem 2

We prove Theorem 2 by an approximation-preserving reduction from the minimum coverage problem (MinC) [4]. Given a nonnegative integer k , a finite set U and a collection $\mathcal{C} \subseteq 2^U$, the MinC problem is to find k subsets $\{S_i, 1 \leq i \leq k\}$ from \mathcal{C} in order to minimize $|\bigcup_{i=1}^k S_i|$.

Given a MinC instance, we can construct a CFM instance as follows. For each element in U , we add $M = k + 1$ nodes with a threshold of 1 in the CFM instance. We call them element-nodes. For each subset in \mathcal{C} , we add a node with a threshold of 1 in the CFM instance. We call them subset-nodes. If an element belongs to a subset in the MinC instance, we produce M edges from the corresponding subset-node to the corresponding M element-nodes. Next, we add a node with a threshold of 0, which is called zero-node. We produce edges from the zero-node to each subset-node. Finally, let the budget be $|\mathcal{C}| - k$ in the CFM instance. It is clear that this reduction can be completed in polynomial time and the maximum threshold equals 1. Denoting the length of the MinC instance as Len, the length of the reduced CFM instance is less than $(M + 1)$ Len. Fig. 2 shows an example of this reduction.

Let OPT_{MinC} denote the optimal value of a MinC instance, and OPT_{CFM} denote the optimal value of the reduced CFM instance.




Fig. 2. An example of the reduction from a MinC instance to a CFM instance with the maximum threshold of 1. In the MinC instance, $U = \{A, B, C, D, E\}$, $\mathcal{C} = \{\{A, B\}, \{B, C\}, \{D\}, \{E\}\}$.

Lemma 5. $M \cdot OPT_{MinC} + k + 1 = OPT_{CFM}$.

Proof. It is clear that any optimal set of edges to delete only contains the edges from the zero-node to subset-nodes. We delete an edge from a zero-node to a subset-node in the CFM instance if and only if the corresponding subset is not chosen in the MinC instance. The optimal solution of the MinC instance minimizes the size of the union of the chosen subsets. Meanwhile, the corresponding set of deleted edges minimizes the number of failed element-nodes. \square

Lemma 6. *For a constant $\epsilon > 0$, if there is a $\frac{1}{2}n^\epsilon$ -approximation algorithm for the CFM problem with the maximum threshold of 1, there exists an $n^{2\epsilon}$ -approximation algorithm for the MinC problem.*

Proof. Given a MinC instance, we can get a CFM instance with the maximum threshold of 1 according to the above reduction. The length of the MinC instance is denoted as Len. The length of the reduced CFM instance is less than $(M + 1)$ Len.

Let E^* be a $\frac{1}{2}n^\epsilon$ -approximation solution of the reduced CFM instance. The budget is $|\mathcal{C}| - k$ in the CFM instance, which is less than or equal to the number of subset-nodes. If E^* contains edges from subset-nodes to element-nodes, we replace these edges with edges from the zero-node to subset-nodes. This procedure will not increase the number of failed nodes. Then, we choose k subsets, the corresponding subset-nodes of which are failed, as an approximate solution of the MinC instance. Let VAL_{CFM} be the approximate value of the reduced CFM instance and VAL_{MinC} be the approximate value of the MinC instance. Because 1 covered element corresponds to M failed element-nodes,

$$M \cdot VAL_{MinC} \leq VAL_{CFM}. \quad (4)$$

According to Eq. (4) and Lemma 5,

$$\begin{aligned} VAL_{MinC} &\leq \frac{VAL_{CFM}}{M} \\ &< \frac{\frac{1}{2}((M + 1) \text{Len})^\epsilon OPT_{CFM}}{M} \\ &< \frac{\frac{1}{2}((M + 1) \text{Len})^\epsilon (M \cdot OPT_{MinC} + k + 1)}{M} \\ &< \text{Len}^{2\epsilon} OPT_{MinC}. \quad \square \end{aligned}$$

Theorem 2 follows from Lemma 6 and the inapproximability result of the MinC problem established by S.A. Vinterbo [4].

References

- [1] A.G. Haldane, R.M. May, Systemic risk in banking ecosystems, *Nature* 469 (2011) 351–355.
- [2] L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, E. Tardos, Which networks are least susceptible to cascading failures?, in: *Proc. of FOCS 2011*, IEEE CS, 2011, pp. 393–402.
- [3] M. Granovetter, Threshold models of collective behavior, *Am. J. Sociol.* 83 (6) (1978) 1420–1443.
- [4] S.A. Vinterbo, A stab at approximating minimum subadditive join, in: *Proc. of WADS 2007*, Springer, Berlin, Heidelberg, 2007, pp. 214–225.
- [5] S. Eubank, V.A. Kumar, M.V. Marathe, A. Srinivasan, N. Wang, Structure of social contact networks and their impact on epidemics, in: *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, vol. 70, 2006, pp. 181–214.
- [6] A. Hayrapetyan, D. Kempe, M. Pál, Z. Svitkina, Unbalanced graph cuts, in: *Proc. of ESA 2005*, Springer, Berlin, Heidelberg, 2005, pp. 191–202.
- [7] J. Aspnes, K. Chang, A. Yampolskiy, Inoculation strategies for victims of viruses and the sum-of-squares partition problem, *J. Comput. Syst. Sci.* 72 (6) (2006) 1077–1093.
- [8] M. Kimura, K. Saito, H. Motoda, Blocking links to minimize contamination spread in a social network, *ACM Trans. Knowl. Discov. Data* 3 (2) (2009) 9:1–9:23.
- [9] D. Kempe, J. Kleinberg, E. Tardos, Maximizing the spread of influence through a social network, in: *Proc. of KDD 2003*, ACM, 2003, pp. 137–146.
- [10] D. Kempe, Structure and dynamics of information in networks, Lecture notes, University of Southern California, Los Angeles. Accessed October 6, 2013, <http://www-rcf.usc.edu/dkempe/teaching/structure-dynamics.pdf>, 2011.