Firefighting as a Game*

Carme Àlvarez, Maria J. Blesa, Hendrik Molter

ALBCOM Research Group - Computer Science Department
Universitat Politècnica de Catalunya - BarcelonaTech
08034 Barcelona, Spain
alvarez@cs.upc.edu, mjblesa@cs.upc.edu, hendrik.molter@gmail.com

Abstract. The Firefighter Problem was proposed in 1995 [16] as a deterministic discrete-time model for the spread (and containment) of a fire. Its applications reach from real fires to the spreading of deseases and the containment of floods. Furthermore, it can be used to model the spread of computer viruses or viral marketing in communication networks.
In this work, we study the problem from a game-theorical perspective. Such a context seems very appropriate when applied to large networks, where entities may act and make decisions based on their own interests, without global coordination.
We model the Firefighter Problem as a strategic game where there is one player for each time step who decides where to place the firefighters. We show that the Price of Anarchy is linear in the general case, but at most 2 for trees. We prove that the quality of the equilibria improves when allowing coalitional cooperation among players. In general, we have that the Price of Anarchy is in $\Theta(\frac{n}{k})$ where k is the coalition size. Furthermore, we show that there are topologies which have a constant Price of Anarchy even when constant sized coalitions are considered.

Keywords: Firefighter Problem; Spreading Models for Networks; Algorithmic Game Theory; Nash Equilibria; Price of Anarchy; Coalitions.

1 Introduction

The Firefighter Problem was introduced by Hartnell [16] as a deterministic discrete-time model for the spread and containment of fire. Since then, it has been subject to a wide variety of research for modeling spreading and containment phenomena like diseases, floods, ideas in social networks and viral marketing.

The Firefighter Problem takes place on an undirected finite graph $G = (V, E)$, where initially fire breaks out at f nodes. In each subsequent time-step, two actions occur: A certain number b of firefighters are placed on non-burning nodes, permanently protecting them from the fire. Then the fire spreads to all non-defended neighbors of the vertices on fire. Since the graph is finite, at some point each vertex is either on fire or saved. Then the process finishes, because

* This work was supported by grants TIN2012-37930 and grant TIN2007-66523 of the Spanish Government, and project 2014-SGR1034 of the Generalitat de Catalunya.

the fire cannot spread any further. There are several different objectives for the problem. Typically, the goal is to save the maximum possible number of nodes. Other objectives include minimizing the number of firefighters (or time-steps) until the spreading stops, or determining whether all vertices in a specified collection can be prevented from burning.

Most research on the Firefighter Problem (also the work in this paper) considers the case $f = b = 1$, which already leads to hard problems. The problem was proved NP-hard for bipartite graphs [20], graphs with degree three [10], cubic graphs [19] and unit disk graphs [14]. However, the problem is polynomial-time solvable for various well-known graph classes, including interval graphs, split graphs, permutation graphs, caterpillars, and P_k-free graphs for fixed k [11, 15, 20, 14]. Furthermore, the problem is $(1 - 1/e)$-approximable on general trees [5], 1.3997-approximable for trees where vertices have at most three children [18], and it is NP-hard to approximate within $n^{(1-\varepsilon)}$ for any $\varepsilon > 0$ [2]. Later results on approximability for several variants of the problem can be found in [2, 4, 7].

Recently, the scientific community has focused on the study of the parameterized complexity of the problem. It was shown to be fixed parameter-tractable w.r.t. combined parameter "pathwidth" and "maximum degree" [6]. Other important results can be found in [8, 3].

Other variants of the Firefighter Problem include the fractional firefighter [13] and the non-constant firefighter [9, 21] (see also Section 8 in [11] for similar problems).

In this work, we study the Firefighter Problem from a game-theorical perspective. Instead of global coordination algorithms, we define a game where the players decide which nodes to protect. Player i chooses where to place the firefighters at time-step i, independently from the other players (one shot game). Since we consider the case of $b = 1$, every player can protect at most one node in his corresponding turn. We can consider different payoffs for the players, the most natural seems to save as many nodes as possible. At each time-step, the fire spreads automatically as described in the original problem.

To the best of our knowledge, the only existing game-theoretical models to similar problems are those referred to as the vaccination problem [2, 12], the spreading of rumors [25] and competitive diffusion [1, 24, 22, 23]. Those models however focus on information spreading on social networks, and thus take into account other inherent aspects of those scenarios, like preferences, reputation, popularity and other personal traits of the users, and relevance or truthfulness of the information. Our proposal is well-suited to model fighting against spreading phenomena in large networks, where the protection strategy for each time-step is decided by one player, independently from the others.

The paper is organized as follows. In Section 2 we define some basic game-theoretical concepts extensively used along the paper. In Section 3 we introduce the game and analyze the quality of its equilibria. Then we explore the behavior on trees. In Section 4 we introduce a solution concept which allows coalitions of players. We show that this improves the Price of Anarchy, explore the computational complexity of finding equilibria and look at graphs with constant cut-width. Finally, conclusions and directions for future work can found in Section 5.

2 Game-Theoretical Definitions

A strategic game $\mathcal{G} = (\mathcal{N}, \mathcal{S}, u)$ is defined by a set of players $\mathcal{N}$, action sets $\mathcal{S}_i$ for each player $i \in \mathcal{N}$ and utilities $u_i : \mathcal{S} \to \mathbb{R}$, where $\mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_{|\mathcal{N}|}$.

Each player i plays an action $s_i \in \mathcal{S}_i$ and his payoff is $u_i(s)$, where $s = (s_1, \ldots, s_{|\mathcal{N}|})$ is the strategy vector or strategy profile of all players. The quality of the outcome of the game when strategy vector s is played is measured by a so-called social welfare function $W(s)$. Furthermore we denote $(s_{-i}, s'_i) = (s_1, \ldots, s'_i, \ldots, s_{|\mathcal{N}|})$, i.e. strategy vector s, where player i changed his strategy from s_i to s'_i.

Nash Equilibrium. A strategy profile s is an Nash equilibrium, if no player can improve his payoff by changing the strategy he played. Let $\mathcal{E} \subseteq \mathcal{S}$ denote the set of all Nash equilibrium strategies. We say that $s \in \mathcal{E}$ if it holds that:

$$\forall i \in \mathcal{N}, \forall s'_i \in \mathcal{S}_i : u_i(s) \geq u_i(s_{-i}, s'_i).$$

Price of Anarchy. The Price of Anarchy (PoA) of a game $\mathcal{G}$ with respect to a social welfare function W is defined as the ratio between the optimal solution and the worst equilibrium.

$$\text{PoA}(\mathcal{G}, W) = \frac{\max_{s \in \mathcal{S}} W(s)}{\min_{s \in \mathcal{E}} W(s)}.$$

Price of Stability. The Price of Stability (PoS) of a game $\mathcal{G}$ with respect to a social welfare function W is defined as the ratio between the optimal solution and the best equilibrium.

$$\text{PoS}(\mathcal{G}, W) = \frac{\max_{s \in \mathcal{S}} W(s)}{\max_{s \in \mathcal{E}} W(s)}.$$

3 The Firefighting Game

The Firefighting Problem takes place on an undirected graph $G = (V, E)$, where fire breaks out at one node, namely $v_0 \in V$, and incinerates all neighboring nodes at every time-step. We call those nodes *burning*. A fixed number b, called the budget, of firefighters can be placed on nodes to permanently protect them from burning. These nodes are called *defended*. If a node never burns because it is defended or cut off from the fire it is called *safe*. All other nodes are called *vulnerable*. We just consider the case of a $b = 1$.

In order to define a firefighting game, we have to define a set of players $\mathcal{N}$, with $\mathcal{N} = \{1, \ldots, n-1\}$ where $n = |V|$, and for every Player $i \in \mathcal{N}$, his strategy set $\mathcal{S}_i$ and his utility function u_i.

Player i decides which nodes to protect at time-step i. His strategy s_i is the subset of nodes he wants to place firefighters, $\mathcal{S}_i$ denotes the set of all possible strategies for player i. Since we only deal with the case of $b = 1$ we overload notation and instead of subsets of size one, we set the strategies to the vertices themselves or the empty set, i.e. $\mathcal{S}_i = V \cup \{\emptyset\}$. This means that players can choose one node or the empty set as a strategy. Let $s = (s_1, \ldots, s_{|\mathcal{N}|})$ denote the strategy profile of all players.

The outcome of the game is a partition of the vertex set into saved and burned nodes. It is defined in the following way. At time-step 0 the only burning node is v_0. At time-step $i > 0$, two events occur: First player i's node is protected if his action is valid w.r.t. to strategy profile s, i.e. it is neither burning nor already defended at the end of time-step $i - 1$. Second, the each node that is burning at time-step $i - 1$ incinerates all its non-defended neighbors. The process stops when the fire cannot spread any further. Let $\text{Safe}(s) \subset V$ be the set of all nodes that are saved when strategy vector s is played. Furthermore, let $\text{Safe}_i(s) = \text{Safe}(s) \setminus \text{Safe}(s_{-i}, \emptyset)$ be the set of nodes that would burn if player i switched his action to the empty set and let $\text{invalid}(s, i)$ denote the event that player i's action is not valid with respect to strategy profile s.

3.1 Utility Functions

We look at two different functions, one modelling a selfish behavior and the other one modelling a non-profitable behavior. As it turns out, the respective games are equivalent.

a) Selfish Firefighters. In this model, firefighters get paid for the nodes they save. We call this game $\mathcal{G}^{(\text{Selfish})}$. Intuitively, if player i makes a valid move other than the empty set, he gets one unit of currency from each node he helped to save. In other words, he gets paid by all nodes that are safe with respect to the played strategy vector, but would not be safe if he would change his strategy to the empty set. Additionally, he will get charged a penalty if he makes an invalid move. Now let us define the utility function formally.

$$u_i^{(\text{Selfish})}(s) = \begin{cases} -c & \text{if } \text{invalid}(s, i), \\ 0 & \text{if } s_i = \emptyset, \\ |\text{Safe}_i(s)| - \varepsilon & \text{otherwise}, \end{cases}$$

with $0 < \varepsilon < 1$ and $c > 0$. We can see that the definition follows the intuition very closely. Subtracting an ε cost for placing a firefighter makes sure that players always prefer to play the empty set over placing a firefighter on a node that is already safe (which would not be an invalid move).

b) Non-Profit Firefighters. Here we assume that the goal of every firefighter is to save as many total nodes as possible, independently of which firefighters actually save more nodes. We call this game $\mathcal{G}^{(\text{Non-Profit})}$. Formally, we define

$$u_i^{(\text{Non-Profit})}(s) = \begin{cases} -c & \text{if invalid}(s,i), \\ |\text{Safe}(s)| & \text{if } s_i = \emptyset, \\ |\text{Safe}(s)| - \varepsilon & \text{otherwise,} \end{cases}$$

with $0 < \varepsilon < 1$ and $c > 0$.

Notice that in an equilibrium, no player plays an invalid move or puts a firefighter on an already safe node. Also, since we have that $0 < \varepsilon < 1$, the cost of placing a firefighter is less than the benefit of saving one node. Because of that, given that a player does not play the empty set, the ε-value does not affect his preferences. Therefore, we will ignore it in the proofs.

Equivalence of Games. Surprisingly, the behavior of selfish firefighters leads to the same equilibria than the behavior of the non-profit firefighters. It can be shown that the games $\mathcal{G}^{(\text{Selfish})}$ and $\mathcal{G}^{(\text{Non-Profit})}$ have the same sets of equilibria (for a proof see Appendix A). This also implies that

$$\text{PoS}(\mathcal{G}^{(\text{Selfish})}, W) = \text{PoS}(\mathcal{G}^{(\text{Non-Profit})}, W)$$

$$\text{PoA}(\mathcal{G}^{(\text{Selfish})}, W) = \text{PoA}(\mathcal{G}^{(\text{Non-Profit})}, W).$$

Therefore we will use the utility function which is more convenient for the proof. Also, we will for now on refer to the game with $\mathcal{G}$, whenever the respective result holds for both versions of the game.

3.2 Quality of Equilibria

Once we have established a game, we can analyze the quality of the equilibria. In order to do this, we have to define a measure of the social benefit. We look at the simple case of the social welfare being the number of the nodes that are saved, i.e. $W(s) = |\text{Safe}(s)|$. It is easy to argue that equilibria always exist, because every optimal solution that does not contain invalid moves is an equilibrium for non-profit firefighter since it maximizes their utility function.

Price of Stability. In the case of non-profit firefighters, every strategy that maximizes the social welfare also maximizes the utility of every player given that he cannot improve his payoff by switching to the empty set. This makes every optimal solution with respect to the social welfare an equilibrium, given no nodes are protected by firefighters that are already save. An optimal solution like that always exists, since any unnecessary firefighters can be removed without decreasing the social welfare. Therefore, we have the PoS is 1. This is independent of the class of graphs we are considering and holds for every solution concept where players maximize their utility function.

Lemma 1. $\text{PoS}(\mathcal{G}, W) = 1.$ □

Price of Anarchy. In contrast to the PoS, the PoA is very high in this model. We first lower bound the PoA and then show that the bound is tight. For the proofs we use the utility functions of the selfish firefighters.

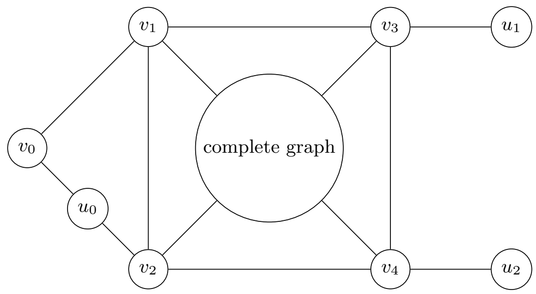

Fig. 1. Family of graphs $G_{PoA}(n) = (V_{PoA}(n), E_{PoA}(n))$. Note that $(v_1, v_4) \in E_{PoA}(n)$ and $(v_2, v_3) \in E_{PoA}(n)$. For better visibility these edges are not drawn in the picture. Further note that $|V_{PoA}(n)| = n$, hence the size of the complete graph is $n-8$ and the nodes of this graph together with nodes v_1, v_2, v_3 and v_4 form a clique of $G_{PoA}(n)$.

Theorem 1. $\text{PoA}(\mathcal{G}, W) \in \Theta(n)$.

Proof. We first prove a lower bound on the PoA, i.e. $\text{PoA}(\mathcal{G}, W) \in \Omega(n)$, and then show that this bound is tight. We look at an instance which has a very bad equilibrium relative to the optimal strategy with respect to the social welfare. Consider the family of graphs $G_{PoA}(n)$ shown in Figure 1.

Recall that the fire starts at v_0. It is easy to see that $s = (\{v_1\}, \{v_2\}, \emptyset^{n-3})$ is the optimal strategy. Only nodes v_0 and u_0 burn, hence the social welfare is $W(s) = n-2$. Furthermore we have that $s' = (\{v_3\}, \{v_4\}, \emptyset^{n-3})$ is an equilibrium. Note that the complete graph is burning after two time-steps, therefore at time-step 3 only u_1 and u_2 are neither burning nor defended. But these nodes are already safe, hence players i with $i > 2$ will not place firefighters on them. Furthermore, players 1 and 2 cannot improve their payoff, since if one of them changes strategy, that player will save at most one node. The social welfare of s' is $W(s') = 4$.

Hence, we have that $\text{PoA}(\mathcal{G}, W) \geq \frac{n-2}{4}$. It follows that $\text{PoA}(\mathcal{G}, W) \in \Omega(n)$. This means that we can only guarantee to save at most constant number of nodes. To argue that this bound is tight, we show that it is always possible to save a constant number of nodes.

By definition player 1 can always place a firefighter on a node before the fire starts spreading. Also any strategy vector s where player 1 plays the empty set is not an equilibrium since he can always save at least one node which cannot be saved by any other player by placing a firefighter to a node adjacent to the original fire. This yields a upper bound of $\text{PoA}(\mathcal{G}, W) \leq n$, and hence $\text{PoA}(\mathcal{G}, W) \in \mathcal{O}(n)$. □

3.3 Price of Anarchy for Trees

Since the PoA is very high in general, let us study the quality of equilibria for particular topologies. Our aim is to prove that there are cases where the quality of the equilibria is close to the quality of an optimal solution. In this section, we look at the PoA on trees. Let $\mathcal{G}_{\text{Tree}}$ denote the Firefighting Game on trees. We show that in contrast to our general result, the PoA is constant for trees. We assume that v_0, the initial fire, is the root of the tree.

Theorem 2. *$PoA(\mathcal{G}_{Tree}, W) \leq 2$.*

Proof. In this proof, we use similar ideas as in the proof of the approximation ratio of a greedy algorithm in a paper by Hartnell and Li [17].

We use the utility functions of the selfish firefighters. This implies that the utility of a player equals the size of the subtree he saves.

Let $\text{opt} = (\text{opt}_1, \dots, \text{opt}_{|\mathcal{N}|})$ be an optimal solution w.r.t to the social welfare, i.e. the optimal action opt_i is the node that is saved at time-step i. Let $s = (s_1, \dots, s_{|\mathcal{N}|})$ be an equilibrium strategy profile of the players. Recall that the optimal actions as well as the player actions are defined as the nodes in the tree that are saved. Let opt_A be the set of optimal actions opt_i, such that there is no player who plays the same action and no player action is an ancestor of opt_i, i.e. $\forall j \in \mathcal{N} : s_j \neq \text{opt}_i \wedge s_j$ is not ancestor of opt_i. Let opt_B denote the remaining optimal actions. Let $P(\text{opt}_i)$ denote the set of action s_j that are successors of opt_i. Let s_A denote the actions of players, that do not have an optimal action as an ancestor, i.e. $\forall j \in \mathcal{N} : \text{opt}_j$ is not ancestor of s_i. Let s_B denote the remaining player actions. Let $\text{save}(a)$ denote the numbers of nodes saved by action a.

Note that in opt_B there are optimal actions where a player plays the same action or a player action is an ancestor. Those corresponding player actions are the ones in s_A. Therefore we have that

$$\sum_{\text{opt}_i \in \text{opt}_B} \text{save}(\text{opt}_i) \leq \sum_{s_i \in s_A} \text{save}(s_i). \tag{1}$$

Because of the equilibrium property, we have that for every $\text{opt}_i \in \text{opt}_A$

$$\text{save}(s_i) \geq \text{save}(\text{opt}_i) - \sum_{s_j \in P(\text{opt}_i)} \text{save}(s_j),$$

because otherwise player i would have an incentive to switch his strategy to opt_i. If we now sum this up over all optimal actions in opt_A, we get

$$\sum_{\text{opt}_i \in \text{opt}_A} \text{save}(\text{opt}_i) \leq \sum_{\text{opt}_i \in \text{opt}_A} \left(\text{save}(s_i) + \sum_{s_j \in P(\text{opt}_i)} \text{save}(s_j) \right).$$

We can split up the sum on the left hand side and get $\sum_{\text{opt}_i \in \text{opt}_A} \text{save}(s_i) + \sum_{\text{opt}_i \in \text{opt}_A} \sum_{s_j \in P(\text{opt}_i)} \text{save}(s_j)$. Note that in the double sum, we sum up exactly over the player actions that have an optimal action as an ancestor i.e. s_B. So we can rewrite this to

$$\sum_{\mathrm{opt}_i \in \mathrm{opt}_A} \mathrm{save}(\mathrm{opt}_i) \leq \sum_{\mathrm{opt}_i \in \mathrm{opt}_A} \mathrm{save}(s_i) + \sum_{s_i \in s_B} \mathrm{save}(s_i).$$

Now we can use Inequality 1 to get

$$\sum_{\mathrm{opt}_i \in \mathrm{opt}} \mathrm{save}(\mathrm{opt}_i) \leq \sum_{\mathrm{opt}_i \in \mathrm{opt}_A} \mathrm{save}(s_i) + \sum_{s_i \in s} \mathrm{save}(s_i).$$

Furthermore, we have that $\sum_{\mathrm{opt}_i \in \mathrm{opt}_A} \mathrm{save}(s_i) \leq \sum_{s_i \in s} \mathrm{save}(s_i)$ which yields

$$\sum_{\mathrm{opt}_i \in \mathrm{opt}} \mathrm{save}(\mathrm{opt}_i) \leq 2 \sum_{s_i \in s} \mathrm{save}(s_i).$$

This shows that an equilibrium strategy saves at least half of the nodes saved by an optimal solution, yielding a PoA of at most 2. □

4 Coalitions

In this section let us consider that players may form coalitions. A coalition is willing to deviate from their strategy as long as no player in the coalition loses payoff and at least one player increases his utility. We show that this affects the PoA. First, we need to introduce a suitable solution concept for coalitions.

We call a strategy vector s an equilibrium strategy with respect to coalition size k, if no set of at most k players can simultaneously change their strategies in such a way that at least one player increases his payoff and no player decreases his payoff. Let $K \subseteq \mathcal{N}$ denote the coalition and s_K a strategy profile of the members of the coalition. We say that coalition K has an *attractive joint deviation* if there is a strategy vector s'_K, such that $u_i(s) \leq u_i(s_{-K}, s'_K)$ for all $i \in K$, and for at least one player in K this inequality is strict.

Let $\mathcal{E}_k \subseteq \mathcal{S}$ denote the set of all equilibrium strategies with respect to coalition size k. We say that $s \in \mathcal{E}_k$, if there is no coalition K of size at most k that has an attractive joint deviation. Formally, we say that $s \in \mathcal{E}_k$ if it holds that:

$$\forall K \subseteq \mathcal{N} \text{ with } |K| \leq k \text{ and } \forall s'_K \neq s_K : s'_K \text{ is not an attractive joint deviation.}$$

Let $\mathcal{G}_k$ denote a firefighting game with coalitions of size at most k. In this case we do not have an equivalence between selfish and non-profit firefighters like in the Nash case. It can be shown that the sets of equilibria of the respective games are different. For a proof of this claim, we refer to Appendix B. From now on we will only consider non-profit firefighters since they resemble the usual objective to save as many nodes as possible.

4.1 Price of Anarchy

Now we analyze the PoA for coalitions and its relation with the coalition size. We can show the following relationship.

Theorem 3. $PoA(\mathcal{G}_k, W) \in \Theta(\frac{n}{k})$.

Proof. To prove this, we first give an upper bound on the PoA for coalition size k. Later we show that this bound is tight. We show the following upper bound. See Appendix C for a detailed proof.

$$\text{PoA}(\mathcal{G}_k, W) \leq \frac{n}{k} - 1.$$

To show this bound, we upper bound the welfare of the optimal solution and lower bound the welfare of the worst equilibrium. Note that if the optimal solution uses k or less time-steps, it can be found by a coalition of size k. Therefore, we assume that in the optimal solution at least in the first $k+1$ time-steps a firefighter is placed on a node. This means that at most $n-k-1$ nodes are saved. We can lower bound the number of nodes saved by the players by k, i.e. the nodes they place firefighters on. This yields a bound of the PoA of at most $\frac{n-k-1}{k} \leq \frac{n}{k} - 1$.

Now we show the following lower bound of the PoA for coalitions of size $k \leq \frac{n-3}{4}$.

$$\text{PoA}(\mathcal{G}_k, W) \geq \frac{n}{k+1} - 3.$$

We construct a family of graphs where the optimal solution saves at least all but $3k+2$ nodes, whereas the worst equilibrium saves at most $k+1$ nodes. Figure 2 shows the construction.

Note that any solution is a lower bound for the optimal solution and every equilibrium is an upper bound for the worst equilibrium in terms of quality.

The solution $s^* = (v_1, v_2, \ldots, v_{k+1}, , \emptyset^{|\mathcal{N}|-k-1})$ saves all but $3k+2$ nodes. This yields a lower bound for the welfare of an optimal solution.

Furthermore, we have that $s = (v'_1, v'_2, v'_3, \ldots, v'_k, \emptyset^{|\mathcal{N}|-k})$ is an equilibrium, since for every joint deviation the players can only save at most k nodes. In this equilibrium they save $k+1$. Now we have a lower bound of the PoA of $\frac{n-3k-2}{k+1} \geq \frac{n}{k+1} - 3$.

Note that this construction uses at least $4k+3$ nodes, hence it is only applicable for coalition sizes up to $k \leq \frac{n-3}{4}$. Since the Price of Anarchy for size $k = \frac{n-3}{4}$ is constant this is no problem for the asymptotic bound.

We have bound the PoA from both sides and it follows that we have the claimed asymptotic behavior. □

It is interesting to see that for linear sized coalitions, we get a constant PoA. For constant coalition sizes however, the PoA is still linear. We can improve this result by fixing a special class of graphs, as we show in the next subsection.

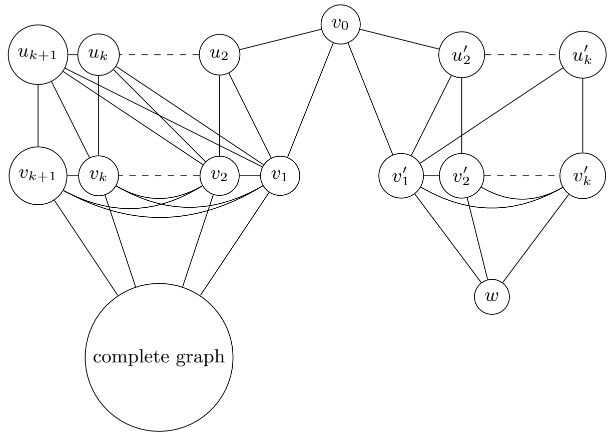

Fig. 2. Family of graphs $G_{PoA}(n,k) = (V_{PoA}(n,k), E_{PoA}(n,k))$, with $|V_{PoA}(n,k)| = n$. Note that the complete subgraph together with nodes v_1 to v_{k+1} form a clique. The nodes v'_1 to v'_k together with w form a clique as well. For every v_i and u_j and for every v'_i and u'_j there are edges (v_i, u_j) and (v'_i, u'_j), respectively, if $i \leq j$. Furthermore, for every u_i and u'_i there is an edge to u_{i+1} and u'_{i+1}, respectively.

4.2 Graphs with constant Cut-width

In this section we explore the impact of the cut-width of a graph on the Price of Anarchy for certain coalition sizes. We make use of results and ideas from Chlebíková and Chopin [6]. In particular, we show that for every family of graphs with constant cut-width there is a constant k, such that the PoA approaches one for coalitions of size k.

The cut-width of a graph G is defined as follows. The Cut-width $\text{cw}(G)$ of a graph G is the smallest integer k such that the vertices of G can be arranged in a linear layout $L = (v_0, \ldots, v_{n-1})$ in such a way that, for every $i \in \{0, \ldots, n-1\}$, there are at most k edges with one endpoint in $\{v_0, \ldots, v_i\}$ and the other in $\{v_{i+1}, \ldots, v_{n-1}\}$. Let $d_L(v_i, v_j) = |j - i|$ denote the distance between two nodes in the linear layout L.

Lemma 2. *If there is one initially burning node, then there exists a protection strategy such that the number of total burned nodes is at most $f(\text{cw}(G))$ for some function $f : \mathbb{N} \to \mathbb{N}$.* □

The proof of a more general version of this claim in contained in the proof of Theorem 2 of [6] and brings us into the position of showing the following lemma.

Lemma 3. *For every family of graphs $G(n) = (V(n), E(n))$ with constant cut-width there is a constant k, such that*

$$\lim_{n \to \infty} \text{PoA}(\mathcal{G}_k, W) = 1.$$

Proof. Let $G(n)$ be a family of graphs with constant cut-width. By Lemma 2 there is a protection strategy s, such that at most $f(\mathrm{cw}(G))$ nodes burn. Now we make use of the fact that the number of time-steps before the spreading of the fire stops is less or equal to the total number of burned vertices. This is because in each time-step at least one node has to burn, otherwise the spreading of the fire would be stopped. Hence we get that with protection strategy s, the fire is contained in at most $f(\mathrm{cw}(G))$ time-steps. Note that we can place at most one firefighter per time-step, therefore a coalition of size $k = f(\mathrm{cw}(G))$ can apply this protection strategy. Furthermore, only a constant number of nodes burn. Hence, asymptotically, we have a PoA of 1. □

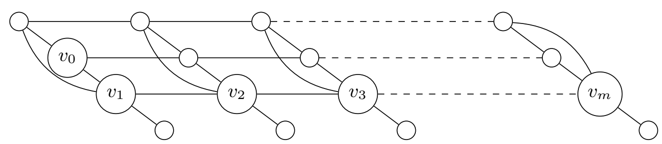

Fig. 3. Family of graphs with constant cut-width.

However, we cannot achieve this without coalitions as the following instance shows. Figure 3 shows a family of graphs. A linear layout is given by the horizontal position of the nodes in the figure. It shows that the cut-width of the graph is at most 6, since every vertical line through the graph crosses at most 6 edges. Without coalitions, saving the nodes v_1 to v_m is an equilibrium, since each player saves one extra node and cannot do better by switching to another node. Note that only a constant fraction of the nodes are saved, whereas in the case of coalition all nodes except a constant number can be saved. This also yields a constant PoA, but one that is asymptotically strictly larger than one.

This shows that for this class of graphs, constant sized coalitions can improve the PoA.

5 Conclusions

We have defined a new strategic game that models the Firefighter Problem. We have shown that in general PoA $\in \Theta(n)$. For trees however, we get a PoA of at most 2, which means that we get equilibria that are close to the optimal solution.

Furthermore, we have shown that the coalition size has a direct effect on the quality of the equilibria. In general we have that PoA $\in \Theta(\frac{n}{k})$, where k is the coalition size. We have shown that there are topologies where PoA approaches 1 for constant sized coalitions, e.g. graphs with constant cut-width.

Note that it is possible to find equilibria in polynomial time for constant sized coalitions. This can be done by best response dynamics. Computing a best response is polynomial since we can try out all possible joint deviations for all possible coalitions of size at most k. With each best response the players impove the total number of saved nodes, hence we converge to an equilibrium in a linear number of iterations. This yields a polynomial time approximation algorithm for the firefighting problem and its approximation ratio equals the PoA of the corresponding game.

We think that the most promising area to explore is the quality of equilibria for other restricted sets of graphs. It is especially interesting to find sets of graphs that have a low PoA for constant sized coalitions.

References

1. N. Alon, M. Feldman, A.D. Procaccia, and M. Tennenholtz. A note on competitive diffusion through social networks. *Information Processing Letters*, 110:221–225, 2010.
2. E. Anshelevich, D. Chakrabarty, A. Hate, and C. Swamy. Approximability of the firefighter problem. *Algorithmica*, 62:520–536, 2012.
3. C. Bazgan, M. Chopin, M. Cygan, M. R. Fellows, F. Fomin, and E. Jan van Leeuwen. Parameterized complexity of firefighting. *Journal of Computer and System Sciences*, 80:1285–1297, 2014.
4. C. Bazgan, M. Chopin, and B. Ries. The firefighter problem with more than one firefighter on trees. *Discrete Applied Mathematics*, 161:899–908, 2013.
5. L. Cai, E. Verbin, and L. Yang. Firefighting on Trees: (1 - 1/e)-Approximation, Fixed Parameter Tractability and a Subexponential Algorithm. *Lecture Notes in Computer Science*, 5369:258–269, 2008.
6. J. Chlebíková and M. Chopin. The firefighter problem: A structural analysis. *Electronic Colloquium on Computational Complexity*, 20:162 (2013)
7. V. Costa, S. Dantas, M. C. Dourado, L. Penso, and D. Rautenbach. More fires and more fighters. *Discrete Applied Mathematics*, 161:2410–2419, 2013.
8. M. Cygan, F. V. Fomin, and E. J. V. Leeuwen. Parameterized Complexity of Firefighting Revisited. *Lecture Notes in Computer Science*, 7112:13–26, 2012.
9. O. N. Feldheim and R. Hod. 3/2 Firefighters Are Not Enough. *Discrete Applied Mathematics*, 161:301–306, 2013.
10. S. Finbow, A. King, G. MacGillivray, and R. Rizzi. The firefighter problem for graphs of maximum degree three. *Discrete Mathematics*, 307:2094–2105, 2007.
11. S. Finbow and G. MacGillivray. The Firefighter Problem: A survey of results, directions and questions. *Australian Journal of Combinatorics*, 43:57–77, 2009.
12. P. Floderus, A. Lingas, and M. Persson. Towards more efficient infection and fire fighting. *International Journal of Foundations of Computer Science*, 24:3–14, 2013.
13. P. Fogarty. Catching the fire on grids. Master's thesis, Department of Mathematics. University of Vermont. USA, 2003.
14. F. V. Fomin, P. Heggernes, and E. J. V. Leeuwen. Making life easier for firefighters. *Lecture Notes in Computer Science*, 7288:177–188, 2012.
15. M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.
16. B. Hartnell. Firefighter! an application of domination. In *25th Manitoba Conference on Combinatorial Mathematics and Computing, University of Manitoba in Winnipeg, Canada*, 1995.
17. B. Hartnell and Q. Li. Firefighting on trees: How bad is the greedy algorithm? *Congressus Numerantium*, 145:187–192, 2000.

18. Y. Iwaikawa, N. Kamiyama, and T. Matsui. Improved Approximation Algorithms for Firefighter Problem on Trees. *IEICE Transactions*, 94-D:196–199, 2011.
19. A. King and G. MacGillivray. The firefighter problem for cubic graphs. *Discrete Mathematics*, 310:614–621, 2010.
20. G. MacGillivray and P. Wang. On the firefighter problem. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 47:83–96, 2003.
21. K. Ng and P. Raff. A generalization of the firefighter problem on Z × Z. *Discrete Applied Mathematics*, 156:730–745, 2008.
22. L. Small and O. Mason. Nash Equilibria for competitive information diffusion on trees. *Information Processing Letters*, 113:217–219, 2013.
23. L. Small and O. Mason. Information diffusion on the iterated local transitivity model of online social networks. *Discrete Applied Mathematics*, 161:1338–1344, 2013.
24. R. Takehara, M. Hachimori, and M. Shigeno. A comment on pure-strategy Nash equilibria in competitive diffusion games. *Information Processing Letters*, 112:59–60, 2012.
25. D. Zinoviev, V. Duong and H. Zhang. A Game Theoretical Approach to Modeling Information Dissemination in Social Networks. *CoRR*, abs/1006.5493, 2010.

A Equivalence of Games

We show that the games $\mathcal{G}^{(\text{Selfish})}$ and $\mathcal{G}^{(\text{Non-Profit})}$ are equivalent in the sense that their sets of equilibria are the same, i.e. $\mathcal{E}^{(\text{Selfish})} = \mathcal{E}^{(\text{Non-Profit})}$.

Lemma 4. $\mathcal{E}^{(\textit{Selfish})} \subseteq \mathcal{E}^{(\textit{Non-Profit})}$.

Proof. We prove this by contradiction. Assume there is a strategy $s \in \mathcal{E}^{(\text{Selfish})}$ which is not an equilibrium for $\mathcal{G}^{(\text{Non-Profit})}$.

If strategy s is not an equilibrium for $\mathcal{G}^{(\text{Non-Profit})}$, this means that there is at least one player i, who can improve the payoff $u_i^{(\text{Non-Profit})}$ by changing his strategy. Let $s' = (s_{-i}, s'_i)$ be the new strategy vector.

If player i can improve $u_i^{(\text{Non-Profit})}$ by playing strategy s'_i, it means more nodes will be saved than when he plays strategy s_i, i.e. $|\text{Safe}(s')| > |\text{Safe}(s)|$.

Let $R = \text{Safe}(s) \setminus \text{Safe}_i(s)$ denote the remaining nodes that are not affected by that change of strategies. We first show that none of the nodes in $\text{Safe}_i(s')$ is contained in R, i.e. $\text{Safe}_i(s') \cap R = \emptyset$. Assume there was a $v \in \text{Safe}_i(s') \cap R$, then it holds that $v \in R$ and $v \in \text{Safe}_i(s')$. The fact that $v \in R$ implies v does not burn if player i would change his strategy to the empty set, meaning that it is not a node that player i helped to save when he plays strategy s'_i, which is a contradiction to $v \in \text{Safe}_i(s')$.

Furthermore we have that $\text{Safe}(s') = R \cup \text{Safe}_i(s')$, i.e. the nodes saved by strategy vector s' are the ones remaining plus the new ones saved by player i when playing s'_i. Since $\text{Safe}_i(s') \cap R = \emptyset$ we have that $|\text{Safe}(s')| = |R| + |\text{Safe}_i(s')|$. Since $R = \text{Safe}(s) \setminus \text{Safe}_i(s)$ and $\text{Safe}_i(s) \subseteq \text{Safe}(s)$, we have that $|R| = |\text{Safe}(s)| - |\text{Safe}_i(s)|$. This yields

$$|\text{Safe}_i(s')| = |\text{Safe}_i(s)| - |\text{Safe}(s)| + |\text{Safe}(s')|. \tag{2}$$

By assumption we have that $|\text{Safe}(s')| > |\text{Safe}(s)|$, therefore we get $|\text{Safe}_i(s')| > |\text{Safe}_i(s)|$ and hence $u_i^{(\text{Selfish})}(s') > u_i^{(\text{Selfish})}(s)$, which is a contradiction to the assumption that $s \in \mathcal{E}^{(\text{Selfish})}$. □

Lemma 5. $\mathcal{E}^{(\textit{Non-Profit})} \subseteq \mathcal{E}^{(\textit{Selfish})}$.

Proof. We prove this by contradiction. Assume there is a strategy $s \in \mathcal{E}^{(\text{Non-Profit})}$ which is not an equilibrium for $\mathcal{G}^{(\text{Selfish})}$.

If strategy s is not an equilibrium for $\mathcal{G}^{(\text{Selfish})}$, this means that there is at least one player i, who can improve the payoff $u_i^{(\text{Selfish})}$ by changing his strategy.

We use Equation 2 from the proof of Lemma 4. This time we have by assumption that $|\text{Safe}_i(s')| > |\text{Safe}_i(s)|$, therefore we get $|\text{Safe}(s')| > |\text{Safe}(s)|$ and hence $u_i^{(\text{Non-Profit})}(s') > u_i^{(\text{Non-Profit})}(s)$, which is a contradiction to the assumption that $s \in \mathcal{E}^{(\text{Non-Profit})}$. □

Corollary 1. $\mathcal{E}^{(\textit{Selfish})} = \mathcal{E}^{(\textit{Non-Profit})}$.

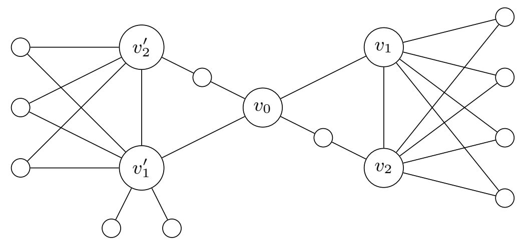

Fig. 4. Example Graph. Note that all nodes have distance at most 2 from the fire.

B Selfish and Non-Profit Firefighters in Coalitions

In the case of coalitions we do not have an equivalence like in the Nash case.

Lemma 6. $\mathcal{E}_k^{(Selfish)} \not\subseteq \mathcal{E}_k^{(Non\text{-}Profit)}$.

Proof. Consider the graph in Figure 4 and assume a coalition size $k \geq 2$. Note that only Players 1 and 2 can make meaningful moves, hence w.l.o.g. we denote the strategy vector with $s = (s_1, s_2)$.

We have that $s = (\{v_1\}, \{v_2\})$ is an equilibrium strategy for the selfish firefighters, since deviating to $s' = (\{v_1'\}, \{v_2'\})$ would decrease the utility of Player 2.

However, s is not an equilibrium strategy for non-profit firefighters, since the joint deviation s' increases the total number of save nodes. $\square$

Lemma 7. $\mathcal{E}_k^{(Non\text{-}Profit)} \not\subseteq \mathcal{E}_k^{(Selfish)}$.

Proof. Consider the graph in Figure 5 and assume a coalition size $k \geq 2$. Note that at most the first 3 players can make meaningful moves, hence w.l.o.g. we denote the strategy vector with $s = (s_1, s_2, s_3)$.

We have that $s = (\{v_1\}, \{v_2\}, \emptyset)$ is an equilibrium strategy for the non-profit firefighters, since deviating to $s' = (\{v_1'\}, \{v_2'\}, \{v_3'\})$ would decrease the total number of save nodes.

However, s is not an equilibrium strategy for selfish firefighters, since the joint deviation $s'_K = (\{v_1'\}, \{v_2'\})$ strictly increases the utility of Player 1 without decreasing the utility of Player 2. $\square$

C Price of Anarchy for Coalitions

Here we analyze the PoA for coalitions and its relation with the coalition size. We can show the following relationship for the firefighting game on arbitrary graphs.

Theorem 2. $PoA(\mathcal{G}_k, W) \in \Theta(\frac{n}{k})$.

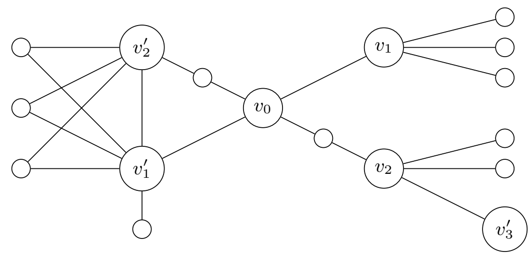

Fig. 5. Example Graph. Note that all nodes have distance at most 3 from the fire.

Proof. To prove this, we first give an upper bound on the PoA for coalition size k. Later we show that this bound is tight. We show the following upper bound.

$$\mathrm{PoA}(\mathcal{G}_k, W) \leq \frac{n}{k} - 1.$$

To show this bound, we upper-bound the welfare of the optimal solution and lower bound the welfare of the worst equilibrium. There are two cases.

Case 1: For instances where there is an optimal solution that uses k or less firefighters, a coalition of size k can always make a joint deviation to that solution. Hence there are no equilibria that have a lower welfare than the optimal solution. It follows that the PoA is 1 for those instances.

Case 2: For instances where every optimal solution uses strictly more than k firefighters, the optimal solution saves at most $n-(k+1)$ nodes. This is because it uses at least $k+1$ time-steps and at least one node burns every time-step, otherwise the fire would be contained. An equilibrium however always saves at least k nodes in this case, because for every strategy profile where the players save less than k nodes, the first k players can jointly deviate to the first k steps of the optimal solution, saving at least the nodes they protect, i.e. k.

Now we have an upper bound for the welfare of the optimal solution and a lower bound of the welfare of the worst equilibrium. This yields a bound of the PoA of at most $\frac{n-k-1}{k} \leq \frac{n}{k} - 1$. Note that the first case is also covered by this bound.

We show the following lower bound of the PoA for coalitions of size $k \leq \frac{n-3}{4}$.

$$\mathrm{PoA}(\mathcal{G}_k, W) \geq \frac{n}{k+1} - 3.$$

Since we have no restriction on the set of underlying graphs, we show that there exist a family of graphs where we get the claimed bound. The optimal solution saves at least all but $3k+1$ nodes, whereas the worst equilibrium saves at most $k+1$ nodes. Figure 6 shows the construction. For a graph of size n, all nodes that are not specifically depicted are inside the complete subgraph. Hence the size of the complete subgraph is $n-4k-3$. Note that the complete subgraph

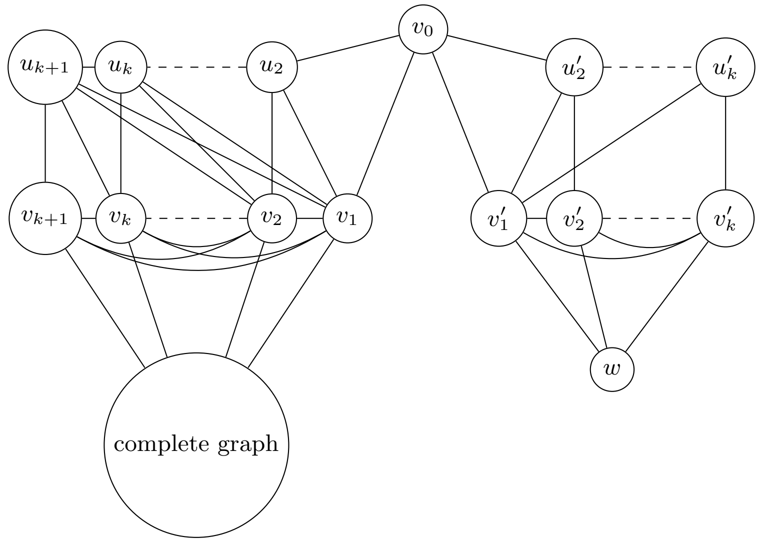

Fig. 6. Family of graphs $G_{PoA}(n,k) = (V_{PoA}(n,k), E_{PoA}(n,k))$, with $|V_{PoA}(n,k)| = n$. Note that the complete subgraph together with nodes v_1 to v_{k+1} form a clique. The nodes v'_1 to v'_k together with w form a clique as well. For every v_i and u_j and for every v'_i and u'_j there are edges (v_i, u_j) and (v'_i, u'_j), respectively, if $i \leq j$. Furthermore, for every u_i and u'_i there is an edge to u_{i+1} and u'_{i+1}, respectively.

together with nodes v_1 to v_{k+1} form a clique. Nodes v'_1 to v'_k and w form a clique as well. We refer to v_0 as the initial fire, nodes v_1 to v_{k+1}, u_2 to u_{k+1} and the complete subgraph as the left part of the graph, and the rest as the right part of the graph.

Note that any solution is a lower bound for the optimal solution and every equilibrium is an upper bound for the worst equilibrium in terms of quality.

The solution $s^* = (v_1, v_2, \ldots, v_{k+1})$ saves all but $3k + 1$ nodes. The nodes that burn are u_2 to u_{k+1} and v_0 as well as the right part of the graph. This yields a lower bound for the welfare of an optimal solution.

Furthermore, we have that $s = (v'_1, v'_2, \ldots, v'_k, \emptyset^{|\mathcal{N}|-k})$ is an equilibrium that saves node w and the nodes that are protected by the firefighters. Since at time-step $k + 1$ there are no vulnerable nodes left, Players i, with $i > k$, have no incentive to deviate from the empty set. Furthermore, we have to argue that for every joint deviation of a coalition of size k at most $k + 1$ nodes can be saved.

We use an inductive argument to show that there is no attractive joint deviation for any coalition of size k into the left part of the graph. Note that for the left hand side of the graph, v_1 is connected to all other vulnerable nodes, which means that if Player 1 does not protect v_1, all vulnerable nodes in the left part will be adjacent to the fire in the next time-step. The next player can protect only one extra node but then everything burns with a total number of two saved nodes. If Player 1 protected v_1, node v_2 assumes the role of v_1 for the next time-step, because it is again connected to all other vulnerable nodes. Hence, for time-step i, with $i \leq k$, we have the following situation. If Players 1 to $i-1$ protect nodes v_1 to v_{i-1}, v_i is connected to all other vulnerable nodes. If Player i does not protect v_i the next player can protect one at most extra node and everything else burns, yielding a number of at most $i+1$ total saved nodes. This implies that in order to save at least $k+1$ nodes in the left part of the graph, Players 1 to k have to protect nodes v_1 to v_k and Player $k+1$ has to play an action different from the empty set. Hence more than k players would have to jointly deviate from strategy profile s.

For the right part of the graph, we can make a symmetric argument, where v'_i assumes the role of v_i for $i \leq k$. This yields that the only way to save at least $k+1$ nodes is to play strategy profile s.

Lastly we have to argue that there is no attractive joint deviation, where nodes from both sides from the graph are protected. Note that both sides of the graph burn in two time-steps if there are no firefighters. This implies that the only way to place firefighters in both parts of the graph is to put one in the left side at the first step and one in the right side at the second, or vice versa. Then at the third time-step, all other nodes in the part where the firefighter was placed in the second time-step are burning. In the part where the first firefighter was placed all unprotected nodes are also burning or adjacent to the fire. This means that at most one extra node can be saved in the third time-step and then the process ends saving at most three nodes.

Now we have shown that $s = (v'_1, v'_2, \ldots, v'_k, \emptyset^{|\mathcal{N}|-k})$ is indeed an equilibrium with social welfare $k+1$.

This yields a lower bound of the PoA of $\frac{n-3k-1}{k+1} \geq \frac{n}{k+1} - 3$, which shows that the bound is asymptotically tight.

Note that this construction uses at least $4k+3$ nodes, hence it is only applicable for coalition sizes up to $k \leq \frac{n-3}{4}$. This is no problem since for this coalition size, the PoA is constant. For larger coalition sizes the PoA can only decrease, since they can make the same joint deviations as smaller ones. This implies that for coalitions of size $k > \frac{n-3}{4}$ the PoA is also constant.

We have bound the PoA from both sides and it follows that we have the claimed asymptotic behavior. $\square$