

Wiretapping a Hidden Network^{*}

Haris Aziz², Oded Lachish¹, Mike Paterson¹, and Rahul Savani³

¹ Department of Computer Science, University of Warwick, CV4 7AL Coventry, UK
{oded,msp}@dcs.warwick.ac.uk

² Institut für Informatik, Universität München, 80538 München, Germany
aziz@tcs.ifi.lmu.de

³ Department of Computer Science, University of Liverpool, L69 3BX Liverpool, UK
rahul.savani@liverpool.ac.uk

Abstract. We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. The *value* of this game is the greatest probability that the wiretapper can secure for hitting the virtual network. The value is shown to be equal the reciprocal of the *strength* of the underlying graph. We provide a polynomial-time algorithm that finds a linear-sized description of the maxmin-polytope, and a characterization of its extreme points. It also provides a succinct representation of all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider. Among these strategies, we efficiently compute the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response. Finally, we show that this unique strategy is the nucleolus of the recently studied simple cooperative *spanning connectivity game*.

Keywords: Network security, nucleolus, wiretapping, zero-sum game.

1 Introduction

We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network, represented by an undirected, unweighted graph. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. A pure strategy of the wiretapper is an edge to tap, and of his opponent, the *hider*, a choice of virtual network, a *connected spanning subgraph*. The *wiretapper* wins, with payoff one, when he picks an edge in the network chosen by the hider, and loses, with payoff zero, otherwise. Thus, the *value* of this game is the greatest probability that the wiretapper can secure for hitting the hidden network. He

^{*} This research was supported in part by EPSRC projects EP/D067170/1, EP/G064679/1, and by the Centre for Discrete Mathematics and its Applications (DIMAP), EPSRC award EP/D063191/1.

does this by playing a maxmin strategy, which is a probability distribution on the edges. The value also equals the smallest probability that the hider can secure, which she does by playing a minmax strategy, which is a probability distribution on connected spanning subgraphs.

Our results. The value is shown to be equal to the *strength* of the underlying graph [6]. We obtain in polynomial time a linear number of simple two-variable inequalities that define the maxmin-polytope, and a characterization of its extreme points. In contrast, the natural description of the maxmin-polytope is as the solutions to a linear program with exponentially many constraints. This allows us to efficiently find all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider, and, among these strategies, the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response. This special maxmin strategy corresponds to the nucleolus of the *spanning connectivity game*, a simple cooperative game [1].

Related work. The strength of an unweighted graph, which has a central role in our work, is also called the edge-toughness, and relates to the classical work of Nash-Williams [8] and Tutte [12]. Cunningham [4] generalized the concept of strength to edge-weighted graphs and proposed a strongly polynomial-time algorithm to compute it. Computing the strength of a graph is a special type of ratio optimization in the field of submodular function minimization [5]. Cunningham used the strength of a graph to address two different one-player optimization problems: the optimal attack and reinforcement of a network. The prime-partition we use is a truncated version of the principal-partition, first introduced by Narayanan [7] and Tomizawa [11]. The principal-partition was used in an extension of Cunningham's work to an online setting [9].

The nucleolus of the spanning connectivity game can be seen as a special maxmin strategy in the wiretap game. The connection between the nucleolus of a cooperative game and equilibrium strategies in a zero-sum game has been investigated before in a general context [10]. However, in many cases the nucleolus is hard to compute. Our positive results for the spanning connectivity game are in contrast to the negative results presented in [1], where it is shown that the problems of computing the Shapley values and Banzhaf values are $\#P$ -complete for the spanning connectivity game.

2 The Wiretap Game

The strategic form of the wiretap game is defined implicitly by the graph $G = (V, E)$. The pure strategies of the wiretapper are the edges E and the pure strategies of the hider are the set of connected spanning subgraphs \mathcal{S} . An element of \mathcal{S} is a set of edges, with a typical element denoted by S . The wiretapper receives payoff one if the edge he chooses is part of the spanning subgraph chosen by the hider, and receives payoff zero otherwise. Thus, the value of the game is

the probability that the wiretapper can secure for wiretapping the connected spanning subgraph chosen by the hider.

Let $\Delta(A)$ be the set of mixed strategies (probability distributions) on a finite set A . By the well-known minmax theorem for finite zero-sum games, the wiretap game $\Gamma(G)$ has a unique *value*, defined by

$$val(\Gamma) = \max_{x \in \Delta(E)} \min_{S \in \mathcal{S}} \sum_{e \in S} x_e = \min_{y \in \Delta(\mathcal{S})} \max_{e \in E} \sum_{\{S \in \mathcal{S}: e \in S\}} y_S . \tag{1}$$

The equilibrium or *maxmin* strategies of the wiretapper are the solutions $\{x \in \Delta(E) \mid \sum_{e \in S} x_e \geq val(\Gamma) \text{ for all } S \in \mathcal{S}\}$ to the following linear program, which has the optimal value $val(\Gamma)$.

$$\begin{aligned} & \max z \\ & \text{s.t. } \sum_{e \in S} x_e \geq z \text{ for all } S \in \mathcal{S} , \\ & x \in \Delta(E) . \end{aligned} \tag{2}$$

Playing any maxmin strategy guarantees the wiretapper a probability of successful wiretapping of at least $val(\Gamma)$. The equilibrium or *minmax* strategies of the hider are $\{y \in \Delta(\mathcal{S}) \mid \sum_{\{S \in \mathcal{S}: e \in S\}} y_S \leq val(\Gamma) \text{ for all } e \in E\}$. Playing any minmax strategy guarantees the hider to suffer a probability of successful wiretapping of no more than $val(\Gamma)$. The following simple observation shows the importance of minimum connected spanning graphs in the analysis of the wiretap game. For a mixed strategy $x \in \Delta(E)$ and pure strategy $S \in \mathcal{S}$, the resulting probability of a successful wiretap is $\sum_{e \in S} x_e$. We denote by G^x the edge-weighted graph comprising the graph G with edge weights $x(e)$ for all $e \in E$. Let $w^*(x)$ be the weight of a minimum connected spanning graph of G^x .

Fact 1. *The set of pure best responses of the hider against the mixed strategy $x \in \Delta(E)$ is*

$$\{S \in \mathcal{S} \mid \sum_{e \in S} x_e = w^*(x)\} .$$

We could define the wiretap game by only allowing the hider to pick spanning trees, however, our definition with connected spanning subgraphs allows a clean connection to the spanning connectivity game.

3 Overview of Results

In this section, we present our results. Proofs of the results appear in [2]. We start with the basic notations and definitions. From here on we fix a connected graph $G = (V, E)$. Unless mentioned explicitly otherwise, any implicit reference to a graph is to G and α is an edge-distribution, which is a probability distribution on the edges E . For ease, we often refer to the weighted graph G^α simply by α , where this usage is unambiguous. For a subgraph H of G , we denote by $\alpha(H)$ the sum $\sum_{e \in E(H)} \alpha(e)$, where $E(H)$ is the edge set of H . We refer to equilibrium strategies of the wiretapper as maxmin-edge-distributions.

Definition 1. For every edge-distribution α , we denote its distinct weights by $x_1^\alpha > \dots > x_m^\alpha \geq 0$ and define $\mathcal{E}(\alpha) = \{E_1^\alpha, \dots, E_m^\alpha\}$ such that $E_i^\alpha = \{e \in E \mid \alpha(e) = x_i^\alpha\}$ for $i = 1, \dots, m$.

Our initial goal is to characterize those partitions $\mathcal{E}(\alpha)$ that can arise from maxmin-edge-distributions α . We start with the following simple setting. Assume that the wiretapper is restricted to choosing a strategy α such that $|\mathcal{E}(\alpha)| = 2$, and $x_2^\alpha = 0$. Thus, the wiretapper’s only freedom is the choice of the set E_1^α . What is his best possible choice? By Fact 1, a best response against α is a minimum connected spanning subgraph H of α . So the wiretapper should choose E_1^α so as to maximize $\alpha(H)$. How can such an E_1^α be found? To answer, we relate the weight of a minimum connected spanning subgraph H of α to E_1^α .

To determine $\alpha(H)$, we may assume about H that for every connected component C of $(V, E \setminus E_1^\alpha)$ we have $E(H) \cap E(C) = E(C)$, since $\alpha(e) = 0$ for every $e \in E(C)$. We can also assume that $|E_1^\alpha \cap E(H)|$ is the number of connected components in $(V, E \setminus E_1^\alpha)$ minus 1, since this is the minimum number of edges in $E(H)$ that a connected spanning subgraph may have. To formalize this we use the following notation.

Definition 2. Let $E' \subseteq E$. We set $C_G(E')$, to be the number of connected components in the graph $G \setminus E'$, where $G \setminus E'$ is a shorthand for $(V, E \setminus E')$. If $E' = \emptyset$ we just write C_G .

Using the above notation, a connected spanning subgraph H is a minimum connected spanning subgraph of α if $|H \cap E_1^\alpha| = C_G(E_1^\alpha) - C_G = C_G(E_1^\alpha) - 1$. Now we can compute $\alpha(H)$. By definition, $x_1^\alpha = \frac{1}{|E_1^\alpha|}$ and $x_2^\alpha = 0$ and therefore

$$\alpha(H) = \frac{C_G(E_1^\alpha) - C_G}{|E_1^\alpha|}.$$

We call this ratio that determines $\alpha(H)$ the cut-rate of E_1^α . Note that it uniquely determines the weight of a minimum connected spanning subgraph of α .

Definition 3. Let $E' \subseteq E$. The cut-rate of E' in G is denoted by $cr_G(E')$ and defined as follows.

$$cr_G(E') := \begin{cases} \frac{C_G(E') - C_G}{|E'|} & \text{if } |V| > 1 \text{ and } |E'| > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

We write $cr(E')$, unless we make a point of referring to a different graph.

Thus, when $|\mathcal{E}(\alpha)| = 2$ and $x_2^\alpha = 0$, a best choice of E_1^α is one for which $cr(E_1^\alpha)$ is maximum. Since E is finite, an E_1^α that maximizes $cr(E_1^\alpha)$ exists.

Definition 4. The cut-rate of G is defined as $opt := \max_{E' \subseteq E} cr(E')$.

By *opt*, we always refer to the cut-rate of the graph G . In case we refer to the cut-rate of some other graph, we add the name of the graph as a subscript.

The value opt is a well known and studied attribute of a graph. It is equal to the reciprocal of the strength of a graph, as defined by Gusfield [6] and named by Cunningham [4]. There exists a combinatorial algorithm for computing the strength, and hence opt , that runs in time polynomial in the *size* of the graph, by which we always mean $|V| + |E|$.

We generalize the above technique to the case that α is not restricted.

Definition 5. For $\ell = 1, \dots, |\mathcal{E}(\alpha)|$ we set

$$cr_\ell^\alpha = \frac{C_G(\cup_{i=1}^\ell E_i^\alpha) - C_G(\cup_{i=1}^{\ell-1} E_i^\alpha)}{|E_\ell^\alpha|}.$$

Proposition 1. Let H be a minimum connected spanning subgraph of α . Then $|E(H) \cap E_\ell^\alpha| = |E_\ell^\alpha| cr_\ell^\alpha$ for every ℓ such that $x_\ell^\alpha > 0$.

Using Proposition 1 we can relate the weight of a minimum connected spanning subgraph of α to the sets of $\mathcal{E}(\alpha)$. This relationship also characterizes the maxmin-edge-distributions, which are the edge-distributions whose minimum connected spanning subgraph weight is the maximum possible.

Theorem 1. Let H be a minimum connected spanning subgraph of α and $m = |\mathcal{E}(\alpha)|$. Then $\alpha(H) \leq opt$ and we have $\alpha(H) = opt$ if and only if

1. $cr_\ell^\alpha = opt$ for $\ell = 1, \dots, m - 1$, and
2. if $cr_m^\alpha \neq opt$ then $x_m^\alpha = 0$.

An immediate implication of Theorem 1 is that opt is an upper bound on the value the wiretapper can achieve. This also follows from the well-known fact that the fractional packing number of spanning trees of a graph is equal to the strength of a graph, which in turn follows from the theorems of Nash-Williams [8] and Tutte [12] on the integral packing number (see also [3]). Since we have already seen that indeed the wiretapper can achieve opt by distributing all probability mass equally over an edge set that has cut-rate opt , we get the following.

Corollary 1. The value of the wiretap game is opt .

We know what the value of the game is and we know a characterization of the $\mathcal{E}(\alpha)$'s for maxmin-edge-distributions α . Yet this characterization does not give us a simple way to find maxmin-edge-distributions. Resolving this is our next goal.

Definition 6. Let $\mathcal{E}_1, \mathcal{E}_2$ be partitions of E . Then \mathcal{E}_1 refines \mathcal{E}_2 if for every set $E' \in \mathcal{E}_1$ there exists a set $E'' \in \mathcal{E}_2$ such that $E' \subseteq E''$.

Thus, there exists a partition of E that is equal to $\mathcal{E}(\beta)$ for some maxmin-edge-distribution β and refines $\mathcal{E}(\gamma)$ for every maxmin-edge-distribution γ . We call such a partition the *prime-partition*. It is unique since there can not be different partitions that refine each other.

Definition 7. The prime-partition \mathcal{P} is the unique partition that is equal to $\mathcal{E}(\beta)$ for some maxmin-edge-distribution β and refines $\mathcal{E}(\gamma)$ for every maxmin-edge-distribution γ .

Theorem 2. The prime-partition exists and can be computed in time polynomial in the size of G .

The prime-partition \mathcal{P} reveals a lot about the structure of the maxmin-edge-distributions. Yet by itself \mathcal{P} does not give us a simple means for generating maxmin-edge-distributions. Using the algorithm for finding \mathcal{P} one can show that, depending on G , there may be a unique element in \mathcal{P} whose edges are assigned 0 by every maxmin-edge-distribution.

Lemma 1. $cr_G(E) \neq opt$ if and only if there exists a unique set $D \in \mathcal{P}$ such that for every maxmin-edge-distribution α and $e \in D$ we have $\alpha(e) = 0$. If D exists then it can be found in time polynomial in the size of G .

From here on we shall always refer to the set D in Lemma 1 as the *degenerate set*. For convenience, if D does not exist then we shall treat both $\{D\}$ and D as the empty set. We use the prime-partition to define a special subset of the minimum connected spanning subgraphs that we call the *omni-connected-spanning-subgraphs*.

Definition 8. A connected spanning subgraph H is an omni-connected-spanning-subgraph if for every $P \in \mathcal{P} \setminus \{D\}$ we have $|E(H) \cap P| = |P| \cdot opt$.

Proposition 2. There exists an omni-connected-spanning-subgraph.

The omni-connected-spanning-subgraphs are the set of the hider's pure strategies that are best responses against every maxmin-edge-distribution.

Proposition 3. For every edge-distribution α such that \mathcal{P} refines $\mathcal{E}(\alpha)$ and $\alpha(e) = 0$ for every $e \in D$ and omni-connected-spanning-subgraph H , we have $\alpha(H) = opt$.

The importance of omni-connected-spanning-subgraphs stems from the following scenario. Assume that \mathcal{P} refines $\mathcal{E}(\alpha)$ and $\alpha(e) = 0$ for every $e \in D$, and let H be an omni-connected-spanning-subgraph. By Proposition 3, we know that $\alpha(H) = opt$. Suppose we can remove from H an edge from $E(H) \cap P$, where P is a nondegenerate element of \mathcal{P} , and add a new edge from another set $P' \setminus E(H)$ in order to get a new connected spanning subgraph. Assume α assigns to the edge removed strictly more weight than it assigns to the edge added. Then the new connected spanning subgraph has weight strictly less than $\alpha(H)$ and hence strictly less than opt , since $\alpha(H) = opt$ by Proposition 3. Consequently, α is not a maxmin-edge-distribution and we can conclude that any edge-distribution β that assigns to each edge in P strictly more weight than to the edges in P' is not a maxmin-edge-distribution. This intuition is captured by the following definition, which leads to the characterization of maxmin-edge-distributions in Theorem 3.

Definition 9. Let $P, P' \in \mathcal{P} \setminus \{D\}$ be distinct. Then P leads to P' if and only if there exists an omni-connected-spanning-subgraph H with $e \in P \setminus E(H)$ and $e' \in P' \cap E(H)$ such that $(H \setminus \{e'\}) \cup \{e\}$ is a connected spanning subgraph.

Definition 10. Let $P, P' \in \mathcal{P} \setminus \{D\}$ be distinct. We say that P is a parent of P' (conversely P' a child of P) if P leads to P' and there is no $P'' \in \mathcal{P}$ such that P leads to P'' and P'' leads to P' . We refer to the relation as the parent-child relation and denote it by \mathcal{O} .

Definition 11. An edge-distribution α agrees with \mathcal{O} if \mathcal{P} refines $\mathcal{E}(\alpha)$ and for every $P \in \mathcal{P} \setminus \{D\}$ that is a parent of $P' \in \mathcal{P} \setminus \{D\}$ and $e \in P, e' \in P'$ we have $\alpha(e) \geq \alpha(e')$, and for every $e \in D$ we have $\alpha(e) = 0$.

Theorem 3. An edge-distribution α is a maxmin-edge-distribution if and only if it agrees with \mathcal{O} .

Theorem 3 defines a linear inequality for each parent and child in the relation \mathcal{O} . Along with the inequalities that define a probability distribution on edges, this gives a small number of two-variable inequalities describing the maxmin-polytope. In [2], we characterize the extreme points of the maxmin-polytope.

Theorem 4. The parent-child relation \mathcal{O} can be computed in time polynomial in the size of G .

The wiretapper will in general have a choice of infinitely many maxmin-edge-distributions. To choose a maxmin-edge-distribution, it is natural to consider refinements of the Nash equilibrium property that are beneficial to the wiretapper if the hider does not play optimally. First we show how to minimize the number of pure best responses of the hider. To do this, we use the relation \mathcal{O} to characterize a special type of maxmin-edge-distribution which achieves this. We call this a prime-edge-distribution. The prime-edge-distributions are characterized by the following lemma.

Definition 12. A maxmin-edge-distribution α is a prime-edge-distribution if the number of the hider’s pure best responses against it is the minimum possible.

Lemma 2. An edge-distribution γ is a prime-edge-distribution if and only if $\gamma(e) > 0$ for every $e \in E \setminus D$, and for every $P, P' \in \mathcal{P} \setminus \{D\}$ such that P is a parent of P' and every $e \in P, e' \in P'$, we have $\gamma(e') > \gamma(e)$.

Using this characterization one can easily check whether α is a prime-edge-distribution and one can also easily construct a prime-edge-distribution.

We have already seen how to minimize the number of pure best responses of the hider, by playing a prime-edge-distribution. We now show how to uniquely maximize the weight of a pure second-best response by choosing between prime-edge-distributions. This maximizes the least punishment that the hider will incur for picking a non-optimal pure strategy.

Against a prime-edge-distribution, the candidates for pure second-best responses are those connected spanning subgraphs that differ from omni-connected-spanning-subgraphs in at most two edges. For each parent and child we have at

least one of these second-best responses. A second-best response either is a best response with one extra edge, or it differs from a best response in two edges, where it has one less edge in a child of \mathcal{O} and one more in the child's parent.

We are only interested in the case that $opt < 1$, since the graph has $opt = 1$ if and only if it contains a bridge, in which case the value of the game is 1 and the hider does not have a second-best response. So we assume that $opt < 1$.

Intuitively, to maximize the weight of a second-best response, we want to minimize the number of distinct weights. The minimum number of distinct positive weights we can achieve for a prime-edge-distribution is equal to the number of elements in the longest chain in the parent-child relation. This motivates the following definition.

Definition 13. We define $\mathcal{L}_1, \mathcal{L}_2, \dots$ inductively as follows. The set \mathcal{L}_1 is all the sinks of \mathcal{O} excluding D . For $j = 2, \dots$, we have that \mathcal{L}_j is the set of all the sinks when all elements of $\{D\} \cup (\cup_{i=1, \dots, j-1} \mathcal{L}_i)$ have been removed from \mathcal{O} .

Note that \mathcal{O} is defined only over nondegenerate elements of \mathcal{P} and hence the degenerate set is not contained in any of $\mathcal{L}_1, \mathcal{L}_2, \dots$.

The following theorem shows that there is a unique prime-edge-distribution that maximizes the difference between the payoff of a best and second-best response. This unique prime-edge-distribution turns out to be the nucleolus of the spanning connectivity game, as explained in [2]. For convenience, we refer to this strategy as the nucleolus.

Theorem 5. Let $L_i = \cup_{E' \in \mathcal{L}_i} E'$ for $i = 1, \dots, t$. Let

$$\kappa = \frac{1}{\sum_{i=1}^t i \cdot |L_i|}.$$

The nucleolus ν has $\nu(e) = i \cdot \kappa$ for every $i \in \{1, \dots, t\}$ and $e \in L_i$ and $\nu(e) = 0$ otherwise.

References

1. Aziz, H., Lachish, O., Paterson, M., Savani, R.: Power indices in spanning connectivity games. In: AAIM: Algorithmic Aspects in Information and Management, pp. 55–67 (2009)
2. Aziz, H., Lachish, O., Paterson, M., Savani, R.: Wiretapping a hidden network. Technical report (2009), <http://arxiv.org/abs/0909.5293>
3. Chakrabarty, D., Mehta, A., Vazirani, V.V.: Design is as easy as optimization. In: ICALP: International Colloquium on Automata, Languages and Programming, pp. 477–488 (2006)
4. Cunningham, W.H.: Optimal attack and reinforcement of a network. J. ACM 32(3), 549–561 (1985)
5. Fujishige, S.: Submodular Functions and Optimization. Annals of Discrete Mathematics, vol. 58. Elsevier, Amsterdam (2005)
6. Gusfield, D.: Connectivity and edge-disjoint spanning trees. Inf. Process. Lett. 16(2), 87–89 (1983)

7. Narayanan, H.: Theory of matroids and network analysis. PhD thesis, IIT, Bombay (1974)
8. Nash-Williams, C.S.A.: Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.* 36, 445–450 (1961)
9. Patkar, S., Narayanan, H.: Fast on-line/off-line algorithms for optimal reinforcement of a network and its connections with principal partition. In: *FST TCS: Foundations of Software Technology and Theoretical Computer Science*, pp. 94–105. Springer, Heidelberg (2000)
10. Potters, J.A.M., Tijs, S.H.: The nucleolus of a matrix game and other nucleoli. *Math. Oper. Res.* 17(1), 164–174 (1992)
11. Tomizawa, N.: Strongly irreducible matroids and principal partition of a matroid into strongly irreducible minors. *Electron. and Commun.* 59(A), 1–10 (1976)
12. Tutte, W.T.: On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.* 36, 221–230 (1961)