

A Framework for Incentive Compatible Topology Control in Non-Cooperative Wireless Multi-Hop Networks

S.Eidenbenz*

P.Santi†

G.Resta†

ABSTRACT

We study the problem of building and maintaining a network topology with certain desirable features in a wireless multi-hop network where nodes behave as selfish agents. We first provide examples showing that existing topology control approaches are not resilient to strategic node behavior, indicating the need to consider possible selfish node behavior at the design stage. To this end, we propose a general framework that can be used as a guideline in the design of incentive compatible topology control protocols. As an example of how to apply our framework to specific topology control protocols, we present an incentive compatible distributed algorithm for building the minimum spanning tree (MST), which is a very well-known topology control approach. We test the economic validity of our protocol through simulation of wireless networks under various advanced signal loss models on the physical layer. To the best of our knowledge, this is the first incentive compatible realization of topology control presented in the literature.

1. INTRODUCTION

Topology formation and maintenance are key tasks for any wireless multi-hop network. In fact, nodes in a wireless multi-hop network typically have the capability of adjusting their transmit power below a maximum value. Since decreasing the transmit power has the positive effect of reducing the interference level in the network and the power consumption at the node, but an excessive reduction of the node transmit power levels could lead to network disconnection, the goal of the topology formation task is to appropriately set the node transmit power levels in such a way that the resulting network topology has the desired features while reducing the interference level and the node power consumption as much as possible. In addition, as nodes can move and dynamically join/leave the network, the task of topology maintenance comes into play, with the goal of reconfiguring the node transmit power levels so that to maintain the desired network topology.

Several protocols for distributed topology formation and maintenance in wireless multi-hop networks, called *topology*

control protocols, have been recently introduced in the literature. Formally, the topology control task (TC for short) can be described as follows. We are given a communication graph $G = (N, E)$, where N is the set of network nodes and E is the set of all possible links in the network, i.e. the set of all wireless links (u, v) that are sustainable when node u transmits at maximum power and the other network nodes (including node v) are silent. For this reason, G is also called the *max-power graph*. The topology control task consists of selecting a subset of the edges in E such that the selected link set satisfies a number of desirable properties such as connectivity, sparseness, planarity, etc. To account for node mobility and dynamic joining/leaving of network nodes, it is typically assumed that the topology control task is executed periodically, or on demand when excessive link failures occur that point to topology changes. Once the desired network topology has been determined, it is assumed that packets are sent through the network along the selected subset of links only.

Although a great variety of topology control concepts and protocols have been proposed in the literature (see, e.g., [5, 6, 19, 23, 25, 27]), all of them are based on the assumption that individual network nodes act in an altruistic, non-selfish way for the common good of having a well-functioning network. Unfortunately, this assumption does not hold in all application scenarios in which nodes are owned by different, independent, profit-maximizing entities, such as an ad hoc phone user, certain types of wireless mesh networks, and so on. As it will be discussed in detail in Section 2, selfish node behavior has a disruptive effect on topology control protocols, since the individual goals of a network node often conflict with the goals of the network designer. As a consequence of this, current approaches to the topology control problem are doomed to perform poorly in a non-cooperative wireless multi-hop network, unless adequate countermeasures are taken.

We believe solving the problem of stimulating cooperation in the task of building and maintaining the network topology is fundamental to a successful realization of the wireless multi-hop paradigm in a non-cooperative environment as this task is a necessary building block on which protocols at higher and lower layers rely. In other words, it is quite unrealistic to assume (as, for instance, has been done in [12]) that network nodes act selfishly in performing the routing task, while they are willing to cooperate in the task of topology control.

In this paper, we introduce a framework to stimulate selfish network nodes to cooperate in the formation and maintenance of the network topology. Our framework is based on game-theoretic concepts from mechanism design, and it can be used as a guideline in the design of incentive compatible topology control protocols for wireless multi-hop networks. As an example of application of our framework, we present an incentive compatible implementation of a popular topology control approach, and we formally prove that it is in the best interest of the network nodes to behave according to the

*Los Alamos National Labs, NM, USA.

†IIT-CNR, Pisa, ITALY




Figure 1: Disruptive effect of selfish node behavior in constructing the MST. Edges are labeled with their weight, which is proportional to the link length. Links in the constructed topology are in bold.

specifications of our proposed incentive compatible topology control protocol. An incentive compatible realization of another popular topology control approach (the k -closest neighbors graph) is presented in the full version of the paper [13]. **Related Work.** While the problem of stimulating cooperation in wireless networks has been addressed at various levels of the network architecture (e.g., at the MAC [8, 21, 26], routing [2, 7, 12, 28, 34, 33], and application layer [4, 14]), the problem of building and maintaining the network topology in a non-cooperative wireless multi-hop network has received comparatively little attention.

A few papers have studied Nash equilibria in topology control-type games.¹ The network model presented in [11] is similar to our network model. However, the analysis reported in [11] is based on a centralized approach to the topology control problem, and does not give any clue on how nodes could be motivated to cooperate in the construction of a certain desired network topology. The network creation game defined and analyzed with respect to Nash equilibria in [15, 1] can be viewed as a topology formation game in spirit similar to our approach. The network creation game does not aim at building a specific topology (such as an MST), but rather aims at optimizing a social choice function. Our method of choice is mechanism design and not only Nash equilibrium analysis, which can be viewed as complementary approaches. A second network creation game and a corresponding mechanism is defined in [29] with a focus on agents joining and leaving the network; again, the focus is not on building a specific topology, but rather optimize a social choice function.

A series of papers [16, 31] have analyzed the problem of sharing the cost of multi-cast trees. This problem is certainly related to the concrete implementation of our framework for MST. The multi-cast nature changes the game considerably, however.

A last related branch of research has focused on studying cooperation issues in the power control problem arising in cellular networks, which consists in varying the transmit power level at the base station and/or at the mobile phone so that to optimize the link quality [20, 30]. The focus in power control is different than in the case of topology control: namely, optimizing a single transmission between a mobile user and the base station, instead of forming and maintaining a com-

munication graph with certain desired features.

2. THE CASE FOR INCENTIVE COMPATIBLE TC

In this section, we present examples of the disruptive effect of selfish node behavior on some of the topology control protocols introduced in the literature. Before presenting the examples, we have to define a model of selfish node in a multi-hop wireless network.

A selfish node aims at increasing the benefit it gets from executing the protocol, while reducing as much as possible the incurred costs in the constructed topology. In the context of topology control, it is reasonable to assume that a node maximizes its benefit when it is connected to as many other nodes as possible (i.e., when the network is connected), and that a node incurs a cost which is proportional to the number and/or power cost of the links incident into it in the selected network topology. A formal definition of utility of a node in the context of topology control is deferred to the next section.

Network connectivity is a global property which cannot be verified locally, so a node u might not be able to check whether a certain power level setting at u results in a globally connected topology. To account for this, we consider two different models of selfish node: *i*) the *global* model, in which a node u somehow knows whether a certain power level setting at u results in a globally connected topology, and *ii*) the *local* model, in which a node u can only verify local properties of the generated network topology (e.g. number and positions of neighbors in the constructed topology). The global model, although possibly unrealistic in many application scenarios, deserves to be investigated since the selfish node is assumed to be ‘more powerful’ than in the local model. In other words, if a certain topology control protocol is shown to be resilient to selfish node behavior in the global model, it retains the same property in the weaker local model.

Let us now turn our attention to some of the topology control approaches introduced in the literature. One of the most studied approaches to the topology control problem is based on the computation of the MST on the maxpower graph, where links are assigned weights according to some criteria (e.g., link length, power cost, expected interference on the link, and so on). The interest in the MST is motivated by the fact that it is the topology of least total cost that maintains the network connected. Example of protocols based on this concept can be found in [6, 19, 23, 25].

¹Quite informally, a set $S = \{S_1, \dots, S_n\}$ of strategies for nodes u_1, \dots, u_n is a Nash Equilibrium if no node u_i has an incentive to deviate from its strategy S_i , provided the other nodes do not change their strategy.

Consider the situation depicted in Figure 1-a), which represents the MST computed on a certain maxpower graph under the assumption that all the nodes act unselfishly. In this situation, the goal (which is altruistically pursued by all network nodes) is to build a connected topology of minimum cost. Note that this globally optimal solution might be suboptimal from a certain node's point of view. For instance, consider node u in Figure 1-a). The cost incurred by u in the MST is 4.8, which equals the cost of the MST edges incident into u . However, if the goal of node u were to maximize its utility, the topology represented in Figure 1-b) would be the optimal choice. In fact, with this network topology node u would still be connected to all the other network nodes (i.e., its benefit would be maximized), but the incurred cost would be minimized (1.3, instead of 4.8).² Note that the topology of Figure 1-b) is suboptimal under a network-wide perspective, since its total cost is 9.3 instead of 8.1.

Thus, we are in a situation in which the designer goal (building a connected topology of minimum cost) is in contrast with the selfish goal of a network node (building a topology which maximizes its own utility). The consequence of selfish node behavior on MST-based topology control is disruptive: a selfish node has no interest in cooperating in the MST construction, while it is instead motivated to build a different, suboptimal network topology. The situation is even worse if several nodes act selfishly, since in this case the constructed network topology in general can be very different from the optimal topology, and it might even not satisfy fundamental properties such as connectivity.

Examples similar to the ones reported above can be easily found for virtually all of the topology control approaches proposed in the literature. This holds true also for local schemes such as KNeigh [5]³, and cone-based TC (CBTC) [27]. Thus, selfish node behavior must be accounted for at the *design stage* of a topology control approach. In the following, we propose a framework for designing TC protocols resilient to selfish node behavior.

3. A GENERAL FRAMEWORK FOR INCENTIVE COMPATIBLE TC

Before introducing our framework, we need to model the topology control task as a game. We recall that the topology control task can be concisely defined as follows. We are given the maxpower graph $G = (N, E)$, where N is the set of network nodes and E is the set of all possible wireless links that can be sustained in the network. Each link $e = (u, v)$ in the maxpower graph is assigned a weight w_e , which is defined according to some criteria (e.g., link length, link power cost, expected interference on the link, and so on). The goal of TC is to select a subset E' of the edge set, such that the induced communication graph satisfies some desirable properties (e.g., connectivity, sparseness, low interference level, and so on), while the cost of E' is reduced as much as possible. The cost of E' can be defined as the sum of the weights of the edges in E' , or as the maximum of the weights of the edges in E' , or in other ways.

In the remainder of this paper we make the following assumptions:

- A1. the edge weight is defined as follows:

$$w_e = n_e \cdot P_{uv}, \quad (1)$$

²Here, we are implicitly assuming the global selfish node model, according to which a node u can somehow verify that the topology represented in Figure 1-b) is connected.

³A specific example based on this TC protocol is presented in the full version of the paper [13].

where P_{uv} represents the minimum transmit power needed to sustain the link $e = (u, v)$ with the desired properties (e.g., at most a given BER at a certain data rate), and n_e denotes the number of packets that will be sent across link e until the next topology control protocol execution.

- A2. the cost function on the selected edge set E' is defined as the sum of the weights of the edges in E' . Formally,

$$c(E') = \sum_{e \in E'} w_e.$$

We remark that assumption A1. above is made for the sake of presentation only, and that the framework proposed in this paper remains valid for arbitrarily defined edge weights.

Parameter n_e in (1) depends on the set of sessions which take place in the time interval between two successive executions of the topology control protocol, and on how these sessions are routed through the network. Unless accurate information about the expected data traffic and the routing algorithm are known, it is difficult to predict the value of n_e for each link in the network. For this reason, in the following we assume all n_e 's in the graph to be equal, and, consequently, we remove n_e from the definition of edge weight. Again, we remark that this assumption is made for the sake of presentation only, and that our framework can be applied with no modification in case of different n_e values on the links.

A final assumption made in our model concerns the wireless medium, which is assumed to be *symmetric*: denoting with P_{uv} the minimum power required to sustain the link from u to v , and with P_{vu} the minimum power required to sustain the reversed link from v to u , we assume $P_{uv} = P_{vu}$.

In order to come up with a game-theoretic model for topology control, we need to determine what the utility functions of our players (i.e., network nodes) are, and what strategies they could follow in order to maximize their gain. Our goal as protocol designers is then to create a mechanism based on monetary transfers that makes behaving in accordance with the prescribed topology control protocol a *dominant strategy*⁴ for the nodes. A topology control protocol (including the rules for the monetary transfer) that satisfies this property is said to be *incentive compatible*⁵.

When defining a utility function, we have to distinguish a case in which a node is content with the outcome of the game (in our setting, the outcome of the game is the selected edge set E'), as opposed to an uncontent case. As discussed in Section 2, the status of being content or not depends on the type of information the node has access to: in the global selfish node model, the player is assumed to be able to verify global network properties, while in the weaker local selfish node model the player can only verify local properties of the network topology. Example of situations in which a node is content can be 'I have a connection to all the other network nodes' (global model), and 'I have a direct, bi-directional link to at least my k closest neighbors' (local model). However, we outline that the framework described here in principle can be applied independently of the property which makes a node content.

To define the utility function of the player, we need also to determine the incurred costs of the player for a given outcome of the game. According to our global cost model, we assume that the incurred cost of node u for the outcome E'

⁴A strategy S for a player p is said to be *dominant* if, no matter what the strategies the other players play, S is the utility-maximizing strategy for p .

⁵The terms *truthful* and *strategy-proof* are also used in the game theory literature.

equals the sum of the cost of the links of E' incident into u . This cost model is coherent with a per-packet approach to topology control, in which nodes can change the transmit power level on a per-packet basis. In this context, the sum of the links of E' incident into u is an estimation of the cost node u incurs for sustaining the links in E' . The other typical topology control approach considered in the literature is the so-called periodical TC approach, in which a node u transmits packets using a certain power $P(u)$ independently of the actual neighbor to which packets are sent, and the value of $P(u)$ is periodically updated to account for topology changes. This approach to the TC problem naturally leads to a node-based, instead of edge-based, definition of topology cost. How to generalize our framework to account for different definitions of incurred cost of the player is subject of ongoing research.

We are now ready to define the utility function used in our topology control game:

DEFINITION 1 (UTILITY FUNCTION). *Let u be a node in $G = (N, E)$, and let $E' \subseteq E$ be the set of links in the topology built (possibly on a subset of the network nodes) at the end of the protocol execution. Let \mathcal{G} be the (global or local) goal that node u is pursuing, and let $T_{\mathcal{G}}^u$ denote the set of topologies that make node u content according to goal \mathcal{G} . The utility of node u for a given constructed topology E' is defined as follows:*

$$U(u, E') = \begin{cases} -\sum_{e=(u,v) \in E'} w_e + \\ \quad -\text{pay}(u, E') + \text{pr}(u, E') & \text{if } (N, E') \in T_{\mathcal{G}}^u \\ -M & \text{otherwise} \end{cases} \quad (2)$$

Note that the above definition of utility function can be restated in the following, more canonical, form: the utility of player u is $M - \sum_{e=(u,v) \in E'} w_e - \text{pay}(u, E') + \text{pr}(u, E')$ if u is satisfied with the constructed network topology, and 0 otherwise. However, for ease of presentation in the following we use the above defined notion of utility function.

If node u is content with the network topology (N, E') , it incurs a cost of w_e for each adjacent edge $e = (u, v) \in E'$, representing the fact that it will have to transmit packets along these links until the topology is updated. Our framework define monetary transfers that we summarize in the utility function as a payment $\text{pay}(u, E')$ that node u will have to make in order to be allowed to participate in the network, as well as a premium payment $\text{pr}(u, E')$ that u will receive from other nodes, which represents the ‘added value’ that u brings to the network. Finally, if node u is not content with network topology (N, E') , it gets a negative utility of $-M$, where M is a large constant representing the fact that the node was not able to achieve its goal. Since we are considering a scenario in which nodes are selfish, but at the same time they are willing to connect to each other, we assume that $M \gg \text{pay}(u, E'), \text{pr}(u, E')$ and of all the edge weights.

The final step in the definition of our topology control game is determining the set of possible strategies for a player. To this purpose, we observe that every topology control protocol relies on nodes determining the weights w_e of the links to their neighbors. In the most general setting, we can thus assume that a node has at least the option of falsely declaring the weights w_e of its adjacent links. Many other possible cheating behaviors of selfish nodes are possible, but they depend on the specific topology control protocol at hand. For this reason, we do not consider them in our general framework, but we carefully discuss them in the following sections, which are devoted to protocol-specific incentive compatible realizations of topology control.

The field of mechanism design offers a standard solution, called a VCG-based mechanism [22], to incentivize all nodes to participate without lying about the weights of their adjacent links. We present this general scheme here and show in subsequent sections how it can be implemented efficiently for a few specific topology control approaches.

Given a communication graph $G = (V, E)$ and a topology control protocol \mathcal{A} , let $E_{\mathcal{A}} = E'$ denote the set of edges that are in the topology constructed by protocol \mathcal{A} , and let $E_{\mathcal{A}}^{-u}$ denote the set of edges that are in the topology constructed by protocol \mathcal{A} when executed on the maxpower graph that does not contain node u (i.e., the graph with node set $N - \{u\}$). Finally, for any link set E'' let $|E''|$ denote the sum of the weights of all edges in E'' , i.e., $|E''| = \sum_{e \in E''} w_e$. Then, we define the premium for node u as follows:

$$\text{pr}(u, E') = |E_{\mathcal{A}}^{-u}| - |E'| + \sum_{e=(u,v) \in E'} w_e^u, \quad (3)$$

where w_e^u is the cost of edge e as declared by node u .

Note that these premiums need to be funded in some way. While there are several ways of doing this, a standard solution is to equally subdivide these costs among all the $|N| = n$ nodes participating in the network (see e.g. Chapter 9 in [18]). Thus, we define the payment function as follow:

$$\text{pay}(u, E') = \frac{\sum_{u \in N} \text{pr}(u, E')}{n}.$$

This completes the definition of our general framework for an incentive-compatible implementation of topology control: a node u in the network receives a premium which equals the ‘added value’ that node u brings to the network (i.e., the difference between the cost of the computed topology without u and the cost of the topology including u), plus the declared cost incurred by u for joining the network. The overall amount of money which must be paid to the nodes is equally subdivided among the network nodes.

Note that the premium that a node receives might exceed, be equal to, or be less than the payment due by the node. However, even if a node ends up the game paying some money to be part of a network which makes it content, this situation is still preferable to not joining the network, as this would drive down the utility function of the node to $-M$ (we recall that M is much larger than the due payment and the edge weights). In other words, if we assume that every node in the network accepts to join the computed network topology T only if T makes it content, it is easy to see that our proposed framework satisfies *individual rationality*: it is rational for any selfish node u to take part in the game, since the outcome of the game can only increase u ’s utility⁶.

Observe that the premium that a node receives depends on the costs incurred by node u in the computed topology *as declared by node u itself*. So, a selfish node might be tempted to falsely declaring these costs, so that to increase its utility. While a formal proof that this cheating behavior does not increase the utility of the node is deferred to the analysis of the protocol-specific incentive compatible TC implementations, we give here the intuition behind this proof.

The premium received by a node as defined in (3) is composed of three terms: the first term is not influenced by node u ’s declaration, since it is the cost of the topology computed on the graph which does not contain u ; the third term can actually be increased by overdeclaring the cost of the edges

⁶Technically speaking, the above statement is true under the assumption that the cost of the few control messages exchanged among nodes to agree on the final topology T is 0. Similar assumptions have been made in the context of incentive-compatible routing for ad hoc networks [2, 12, 34].

incident into u . However, if these costs are overdeclared by, say, $\Delta > 0$ overall, also the cost of the computed topology $|E_{\mathcal{A}}|$ is increased of the same amount, and the premium due to node u is decreased by Δ . So, node u has no way of increasing its utility (we remark that the two other terms in the utility function of u do not depend on the declared edge weights) by falsely reporting its edge weights.

The framework described in this section is based on the well-known VCG mechanism, which has the nice feature of motivating nodes not to lie when reporting their weights. However, this comes at the price of paying nodes in excess to their real cost for being part of the network (in fact, $|E_{\mathcal{A}}^{-u}| - |E_{\mathcal{A}}|$ is in general a positive quantity), disclosing an economic inefficiency of the mechanism: the sum of the payments due to the nodes (which equals the sum of the premiums paid to the nodes) is in general higher than the cost of the generated topology E' . Unfortunately, it has been shown that, under realistic assumptions, there is no way of removing this inefficiency if the goal is to design an incentive compatible mechanism [22].

A final observation regards the interaction between topology control and routing. Once a desired network topology has been built in an incentive-compatible way, messages circulate in the network, and selfish nodes must be motivated to route packets to destination. In principle, this can be achieved by using one of the incentive-compatible routing protocols for wireless multi-hop networks proposed in the literature [2, 12, 34]. However, topology construction and routing are not independent tasks, since the choice of the topology have an impact on the number and cost of the available source/destination paths. Thus, a global solution that accounts for both topology construction and routing is needed in order to ensure incentive-compatibility. We believe this is a very interesting topic for future research.

4. INCENTIVE COMPATIBLE MST TOPOLOGY CONTROL

4.1 Problem definition and model

In this section we consider the problem of building and maintaining the MST of the maxpower graph G in a scenario in which the network nodes are selfish agents. To simplify the presentation of our protocol, we assume that no two edges e_1, e_2 in G exist such that $w_{e_1} = w_{e_2}$. Note that this assumption can be accomplished by ordering node IDs in lexicographical order, and by breaking ties according to the IDs of the endpoints of the edge. Under this assumption, the MST of G is uniquely defined.

The utility function is the same for all the network nodes, and it is defined as follows:

DEFINITION 2 (UTILITY FUNCTION FOR MST). *Let u be a node in $G = (N, E)$, and let $T = (N_T, E_T)$ be the topology built (possibly on a subset of the network nodes) at the end of the MST protocol execution. The utility of node u is defined as follows:*

$$U(u, T) = \begin{cases} -\sum_{e=(u,v) \in E_T} \frac{w_e}{2} - \\ \quad -\text{pay}(u, T) + \text{pr}(u, T) & \text{if } u \in N_T \\ -M & \text{otherwise} \end{cases} \quad (4)$$

In the above definition we are assuming the global selfish node model. This is consistent with the nature of the MST, whose computation requires global knowledge. Differently from the general framework, we are also assuming that the cost of an edge is equally subdivided between its endpoints.

The goal of each node participating in the protocol execution is to maximize its utility function (selfish agent). To

achieve its goal, a node can deviate from the behavior prescribed by the protocol, for instance by not sending a message, or by reporting false information, or by sending a message more than once. However, nodes are not allowed to coordinate their cheating behaviors in order to form a coalition (*no collusion*)⁷. Our goal, as the protocol designer, is to devise a mechanism (more specifically, a pricing rule) such that *every* node participating in the protocol maximizes its utility when a certain *social optimum* is achieved. This property is known as incentive compatibility in the game theory literature. In simple words, incentive compatibility ensures that the combination of the agents' selfish behaviors results in a desirable "social" behavior.

Returning to the problem at hand, our designer goal is to devise a protocol such that the function

$$\text{Soc}(T) = \sum_{u \in N} U(u, T),$$

which represents the *social utility function*, is maximized when the constructed topology T is the MST built on G .

4.2 The IC-MST protocol

Our incentive compatible protocol for building the MST, called IC-MST, is essentially an incentive compatible implementation of the Prim's algorithm for building the MST, and it is defined as follows.

In order to simplify the presentation, we assume that there exists a network node (the *initiator*) that initiates the protocol at a certain time, and that the initiator node is unique (i.e., no other node in the network can initiate IC-MST execution). This assumption is reasonable, for instance, in wireless Mesh networks, where the initiator is a wireless Access Point, and the goal is to establishing (possibly multi-hop) wireless connections with mobile nodes in the vicinity of the AP. In this scenario, building a MST rooted at the AP is a reasonable choice.⁸ Another assumption needed for the correct function of IC-MST is that every node has at least two neighbors in the maxpower communication graph.

IC-MST proceeds in rounds, adding a new node (with a corresponding edge) to the constructed topology at each round, until all the network nodes are connected ($n-1$ rounds in total). At each round, nodes are requested to pay an amount of money, which is delivered in a secure way to the initiator. The initiator also collects its own money. After round $n-1$, there is a last round called the *premium round*, during which the initiator delivers (in a secure way) the premiums to all the network nodes (including itself).

At round 1, the initiator starts the protocol execution, finding the edge e_1 of minimum cost incident into it, and the second best edge e_1^2 . The node at the other endpoint of edge e_1 joins the network topology, forming a network with 2 nodes (denoted T_1). Edge e_1 is stored in the first element of the array *Edges*[], which keeps track of the network edges (this array and the arrays for payment tracking are stored at the initiator node). The initiator then collects the payments, which amount to $w_{e_1^2}$. The money due is equally shared between the two nodes in T_1 . Furthermore, $w_{e_1^2}$ is stored in the first element of the array *Payments*[], which keeps track of the payments performed in the various rounds.

⁷This is a standard assumption in mechanism design.

⁸We are well aware that the introduction of a central role (in our case the initiator) is against Puritan ad hoc networking principles, however, we point out that in particular in the area of selfish wireless networking (as well as in some others), such a figure has become accepted as a necessary evil (see for instance the central bank concepts of [2, 33, 12]). In addition, many applications naturally offer a central node.

At the generic round i , the initiator asks each node in the current tree T_{i-1} to report the weights of the edges incident into it whose other endpoint is not in T_{i-1} . If no such edge exists, we are done (all the n nodes are included in the topology), and the initiator starts the premium round (see below). If a node u is adjacent to $h > 2$ edges with nodes not in T_{i-1} , it is sufficient that it reports to the initiator the cost of the two edges of minimum cost. If these costs were already reported by u to the initiator in a previous round, node u simply does not report any cost. After all the costs have been collected, the initiator selects the minimum cost edge $e_i = (u, v)$ as the new edge to be added to T_{i-1} . Edge e_i is stored in $Edge[i]$, and the node at the endpoint of e_i which was not in T_{i-1} (say, node v) is added to the newly formed topology T_i . Before computing the payments, the initiator sends a message to node v , asking him to send the list of its neighbors, along with the corresponding link costs. This message, which is necessary to prevent some types of cheating behavior (see the proof of Theorem 1), is encrypted and digitally signed by node v , so that the initiator can rely on this information. After topology T_i has been built, the initiator computes the payments due by the nodes in T_i (including itself), according to the current rule:

$$\begin{aligned} Payments[i] &= Payments[i-1] + w_{e_i^2} \\ pay(v) &= \frac{Payments[i]}{i+1} \\ pay(u) &= \frac{Payments[i]}{i+1} - \frac{Payments[i-1]}{i} \end{aligned}$$

for each $u \in T_{i-1}$,

where $Payments[i-1]$ denotes the total payments collected up to round $i-1$, $Payments[i]$ denotes the total payments collected up to round i , and $w_{e_i^2}$ is the cost of the second best edge connecting a node in T_{i-1} with a node outside T_{i-1} .

The above defined payment rule is inspired by the general framework introduced in Section 3. In principle, we want to ensure that all the nodes joining the network equally share the cost of setting up the network topology. This rule should apply independently of the round in which a node joins the network, since otherwise nodes would be motivated to either anticipate or delay the moment at which they join the network. This explains why we keep track of the total payments collected up to round $i-1$: since these have been equally shared amongst the nodes in T_{i-1} , when a new node joins the network the additional payment due by a node in T_{i-1} equals $\frac{Payments[i]}{i+1}$ (the total amount of money that it must pay), decreased by $\frac{Payments[i-1]}{i}$ (the amount of money already paid by the node). On the contrary, the new node joining the network is charged the entire due payment $\frac{Payments[i]}{i+1}$.

Note that at each round the total cost is increased by the cost of the second best edge joining a node in T_{i-1} with a node outside T_{i-1} . Again, this rule is inspired by the general framework for incentive compatible topology control described in Section 3.

Finally, we have to define the premium round, which is again dictated by our general framework. The initiator scans the array $Edges[]$ starting from the first element. For each edge $e_i = (u, v)$, it computes the premiums due to nodes as follows:

$$pr(u) = pr(v) = \frac{Payments[i] - Payments[i-1]}{2}.$$

Note that $Payments[i] - Payments[i-1]$ corresponds to the cost of the second best edge joining a node in T_{i-1} with a node outside T_{i-1} . This premium is equally divided between the nodes at the endpoints of the edge. The premiums due to

the nodes are summed up as new edges are considered, until all the edges in the network topology have been considered. Then, the initiator delivers to each node (including itself) the corresponding premium in a secure way. An example of IC-MST's execution is reported in the full version of the paper [13].

An important point to discuss is how the edge weights are computed in IC-MST, since improperly implementing this task might impair the incentive compatibility property of our protocol.

The idea is to force each node to either overdeclare, or underdeclare, or correctly declare the cost of all the edges incident into it. In other words, we want to avoid that a node u can, say, correctly declare the cost of a certain edge (u, v) , while at the same time overdeclaring the cost of another edge (u, w) . To this purpose, edge weights are computed by exchanging *hello* messages between the nodes: when the initiator starts the protocol execution, it sends a hello message at maximum transmit power, indicating that it is starting the construction of the network topology. The message includes the initiator's ID and the transmit power used to send the message. All the nodes within its radio coverage area⁹, after waiting for a random time (this is to avoid collisions), reply with another hello message sent at maximum power, reporting the ID and transmit power of the sender. When the initiator receives a hello message from a neighbor node, say v , it compares the transmit power P_v included in the message with the received power P_v^R . The difference between these power levels, namely $P_v - P_v^R$, corresponds to the path loss experienced by the wireless link connecting the initiator with v , and the minimum power needed to sustain a link with the desired features to v can be computed accordingly¹⁰. If the required power to sustain the link to v exceeds the initiator's maximum transmit power¹¹, then node v is not included in the initiator's neighbor list, as the link to v is unidirectional and only bi-directional links are considered in our approach.

Note that receiving a hello message for the first time triggers the recipient node to send a hello message in turn (with a random delay). After all the nodes in the network have sent their hello message, all the edge weights in the communication graph can be correctly computed.

At any time during IC-MST execution, a node which does not behave according to the protocol specifications can be identified as a cheater and excluded from the network topology. For instance, a node which tries to send several hello messages to induce false edge weights at its neighbors can easily be identified as a cheater by the nodes in its vicinity, and excluded from the network topology. This repudiation mechanism is fundamental to ensure the incentive compatibility property of IC-MST (see proof of Theorem 1 below). We remark that the necessity of both a VCG-based mechanism and a repudiation mechanism is a unique feature of selfish distributed computation: a repudiation mechanism by itself is typically not provably correct, but works in a distributed setting; a VCG-mechanism by itself is provably correct, but typically only works in a centralized setting.

4.3 Protocol Analysis

⁹We recall that we are not assuming that the coverage area is a perfect circle, nor that all the nodes have the same maximum transmitting range.

¹⁰Given our assumption of symmetric wireless medium, both the endpoints of a link l compute the same value of the minimum power needed to sustain l .

¹¹In general this is possible, since nodes can have different maximum transmit powers.

THEOREM 1. *Assume the maxpower graph G is strongly connected and has minimum node degree $\delta \geq 2$. Then IC-MST is incentive compatible, i.e. a network node maximizes its own utility when it behaves according to IC-MST's specifications.*

PROOF. See Appendix. \square

Note that the incentive compatibility property of IC-MST relies on the fact that payments/premiums can be gathered/delivered to the nodes in a secure way. Indeed, how to implement secure crediting in wireless multi-hop networks is a research field in itself (see, for instance, [33]). While a detailed discussion of this issue is beyond the scope of this paper, we sketch here a repudiation-based mechanism that can be used to motivate nodes to propagate the payment info to/from other nodes in the tree. We have two critical points in IC-MST: *i*) when a new edge $e = (u, v)$ is added at a certain step, nodes u and v must deliver payments to the initiator; and *ii*) the delivery of premiums during the premium round. Regarding *i*), we first notice that, since for every edge in the MST incident into it a node always receives a premium at least equal to the payment for the edge, node u and v will to send the payments to the initiator. The payment message is encrypted, so that intermediate nodes in the path P to the initiator can only forward or drop the message, but they cannot modify it. If one of the nodes in P drops the message, the initiator realizes that something is wrong (it is in fact expecting the payments for edge e), and it can exclude all the nodes in P from the network topology. Hence, an intermediate node is motivated to forward the payment to avoid being excluded from the network topology. The argument for *ii*) is similar: At the end of the premium round, the initiator individually asks to each node (using digitally signed messages) if it has received the premium: if a node u responds 'no' (or does not respond at all), all the nodes in the path P from the initiator to node u are excluded from the network. The use of digital signatures for the reply message prevent nodes in P from forging u 's reply.

THEOREM 2. *The social optimum is achieved when the topology T computed by IC-MST is the MST of the maxpower graph G .*

PROOF. See Appendix. \square

THEOREM 3. *The IC-MST protocol has $O(n^2)$ message complexity.*

PROOF. See Appendix. \square

5. SIMULATIONS

In Section 3, we have observed that the design of an incentive compatible mechanism for network topology formation comes at the expense of an economic inefficiency, i.e., the total of the payments paid by the nodes to join the network exceeds the actual cost of maintaining the formed network topology. This inefficiency has been referred to as the *cost of cooperation* in the literature. In this Section, we evaluate the cost of cooperation of our proposed protocol through simulation.

We have distributed uniformly at random a number n of nodes varying from 20 to 150 in a square area with side equal to 500m. In each experiment, one of the nodes was randomly selected as the initiator of the protocol. We have repeated the experiments selecting the node which is closest to the center

of the deployment region as the initiator, obtaining virtually the same results. The cost of the edges in the maxpower graph was computed according to four different metrics:

- *Free Space*: an edge $e = (u, v)$ is inserted in the maxpower graph if and only if $dist(u, v) \leq r_M$, where r_M is the maximum transmitting range of the nodes. The weight of edge e is $w_e = (dist(u, v))^2$.
- *Log-Normal Shadowing*: In this model, the transmitted signal attenuation at a certain distance is determined by the sum of a deterministic and a random component. The deterministic component gives the average value of the received signal, which is determined by the distance between sender and receiver and by the path loss exponent (set to 2 in our experiments). The random component has log-normal distribution (normal distribution when measured in dBs) with standard deviation σ ($\sigma = 4$ in our experiments).
- *Free Space Level-based*: In practice, the energy consumption of a wireless NIC does not only depend on the power used to transmit the signal, but also on some constant energy consumption due to powering up other circuitry. Furthermore, only a few transmit power levels P_1, \dots, P_{MAX} can be used to send packets. To account for these aspects, we have considered a level-based version of the free space propagation model, in which the rule used to assign weights to the edges is as follows. Let P_i be the minimum power level needed to sustain the link between nodes u and v . The weight of edge $e = (u, v)$ is defined as:

$$w_e = E_{MAX} - (E_{MAX} - E_1) \left(\frac{r_M^2 - r_i^2}{r_M^2 - r_1^2} \right), \quad (5)$$

where E_1 is the power consumption of the wireless NIC when transmitting at power level P_1 , E_{MAX} is the power consumption when transmitting at power level P_{MAX} , and r_i (respectively, r_1) is the transmitting range at power level P_i (respectively, P_1). Observe that we have $w_e = E_1$ if the minimum power level needed to sustain the link is P_1 , and $w_e = E_{MAX}$ if the minimum power needed to sustain the link is P_{MAX} .

- *Log-Normal Shadowing Level-based*: the same as Free Space Level-based, where the minimum power needed to sustain a certain link is computed according to the Log-Normal Shadowing model.

The results of our simulations, averaged over 10000 simulation runs for every value of n , are reported in Figure 2. The cost of cooperation of the IC-MST protocol is obtained by dividing the sum of the node payments by the cost of the MST. Note that the obtained ratio is always above 1, and that smaller values of the cost of cooperation mean better economic efficiency. The results of our simulations show that: *i*) the cost of cooperation of the IC-MST protocol shows a decreasing trend with n , i.e. it remains confined within a reasonable value (below 1.65) even for networks of large size; this implies that the IC-MST approach for topology formation is *feasible from an economic point of view*; *ii*) the cost of cooperation with log-normal shadowing is lower than in the case of free space propagation; and *iii*) the cost of cooperation is considerably reduced (it actually becomes negligible – below 1.04) if the more realistic level-based metrics are used. This is because, with a limited choice of different power levels, it is more likely that the cost of the 'best' edge to be added in the topology and the cost of the 'second best' edge are close to each other. Concluding, *ii*) and *iii*) indicate that

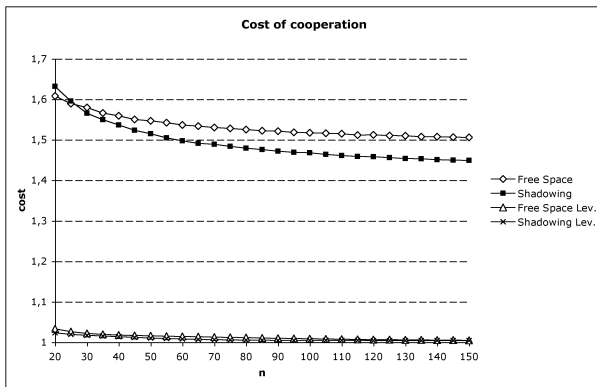


Figure 2: The cost of cooperation of the IC-MST protocol with various network sizes.

IC-MST has the potential to work well in practical scenarios.

6. CONCLUSIONS

Recent technology trends indicate that wireless multi-hop networks formed by nodes belonging to different authorities and governed by conflicting interests will become widespread. This is the case, for instance, of many application scenarios of wireless Mesh networks, and of the general vision of a future in which ubiquitous mobile computing will become reality. In view of these trends, we believe the study reported in this paper constitutes an important building block of next generation wireless multi-hop networks.

The implementations of incentive compatible TC presented in this paper are only examples of application of our framework, that can be applied to other topology control protocols such as, for instance, K-Neigh [5] (as reported in the full version of the paper [13]) and cone-based topology control (CBTC, [27]). Furthermore, the general framework for incentive compatible topology control presented in this paper leaves space for several generalizations, such as accounting for different cost metrics (e.g., max instead of total edge cost). Furthermore, we can adopt a more general notion of benefit that a node gets from a certain network topology T : instead of being either 'content' or 'discontent' with T , a node might display a certain degree d of satisfaction with T , where d is a constant such that $0 \leq d \leq 1$. Whether our proposed framework can be applied with this more general notion of node benefit and with different cost metrics is an open problem, which is matter of ongoing research.

7. REFERENCES

- [1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour and L. Roditty, "On Nash equilibria for a network creation game", *Proc. 17th ACM-SIAM SODA*, pp. 89-98, 2006.
- [2] L. Anderegg, S. Eidenbenz, "Ad hoc-VCG: A Truthful and Cost-Efficient Routing Protocol for Mobile Ad hoc Networks with Selfish Agents", *Proc. ACM Mobicom*, pp. 245-259, 2003.
- [3] N. Ben Salem, J. Hubaux, M. Jakobsson, "Reputation-based Wi-Fi Deployment - Protocols and Security Analysis", *Proc. WMASH '04*, pp. 29-40, 2004.
- [4] D. M. Blough, M. Leoncini, G. Resta, and P. Santi, "The k -NEIGH Protocol for Symmetric Topology Control in Ad Hoc Networks", in *Proc. ACM MobiHoc '03*, pp. 141-152, June 2003.
- [5] M. Burkhart, P. Von Rickenbach, R. Wattenhofer, A. Zollinger, "Does Topology Control Reduce Interference?", *Proc. ACM MobiHoc*, pp. 9-19, 2004.
- [6] L. Buttyan, J. Hubaux, "Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks", *ACM/Kluwer Mobile Networks and Applications*, Vol. 8, No. 5, October 2003.
- [7] M. Cagalj, S. Ganeriwal, I. Aad, J. Hubaux, "On selfish behavior in CSMA/CA networks", *Proc. IEEE Infocom '05*, Miami - FL, USA, March 13-17, 2005.
- [8] J. Ebert, S. Aier, G. Kofahl, A. Becker, B. Burns, A. Wolisz, "Measurements and Simulation of the Energy Consumption of a WLAN Interface," *Tech. Rep. TKN-02-010*, Tech. Univ. of Berlin, June 2002.
- [9] *Cisco Aironet 350 data sheets*, available at <http://www.cisco.com/en/US/products/hw/wireless>.
- [10] S. Eidenbenz, V.S. Kumar, S. Züst, "Equilibria in Topology Control Games for Ad Hoc Networks", *Proc. ACM DIALM-POMC 2003*, pp. 2-11, 2003.
- [11] S. Eidenbenz, G. Resta, P. Santi, "COMMIT: A Sender-Centric Truthful and Energy-Efficient Routing Protocol for Ad Hoc Networks with Selfish Nodes", *Proc. IEEE WMAN*, 2005.
- [12] S. Eidenbenz, P. Santi, "A Framework for Incentive Compatible Topology Control in Non-Cooperative Wireless Multi-Hop Networks", *Tech. Rep. IIT-TR12/2005*, Istituto di Informatica e Telematica, Pisa, Italy, Dec. 2005.
- [13] E. Efstathiou, G. Polyzos, "Fully Self-Organized Fair Peering of Wireless Hotspots", *Wiley European Trans. on Telecommunications*, Vol. 16, n. 5, pp. 471-482, 2005.
- [14] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, Scott Shenker "On a network creation game", *Proc. ACM PODC 2003*, pp. 347 - 351, 2003.
- [15] J. Feigenbaum, C. Papadimitriou, S. Shenker, "Sharing the Cost of Multicast Transmissions", *Journal of Computer and System Sciences*, 63, no. 1, 2001.
- [16] J.R. Green, J.-J. Laffont, "Incentives in Public Decision Making", *North Holland Publishing Company*, 1979.
- [17] N. Li, J. Hou, L. Sha, "Design and Analysis of an MST-based Topology Control Algorithm", *Proc. IEEE Infocom '03*, pp. 2003.
- [18] A.B. MacKenzie, S. Wicker, "Game Theory and the Design of Self-Configuring, Adaptive Wireless Networks", *IEEE Communications Magazine*, Vol. 39, n. 11, pp. 126-131, 2001.
- [19] A.B. MacKenzie, S. Wicker, "Stability of Multipacket Slotted Aloha with Selfish Users and Perfect Information", *Proc. IEEE Infocom '03*, pp. 1583-1590, 2003.
- [20] A. Mas-Colell, M. Whinston, J. Green, *Microeconomic Theory*, Oxford University Press, New York, 1995.
- [21] K. Moaveni-Nejad, X.Y. Li, "Low-Interference Topology Control for Wireless Ad Hoc Networks", *Ad Hoc and Sensor Networks: an International Journal*, Vol. 1, n. 1-2, pp. 41-64, 2005.
- [22] D. Parkes, J. Shneidman, "Distributed Implementations of Vickrey-Clarke-Groves Mechanisms", *Proceedings of ACM AAMAS'04*, 2004.
- [23] R. Ramanathan, R. Rosales-Hain, "Topology Control of Multihop Wireless Networks using Transmit Power Adjustment", *Proc. IEEE Infocom 2000*, pp. 404-413, 2000.

- [24] M. Raya, J. Hubaux, I. Aad, “DOMINO: A System to Detect Greedy Behavior in IEEE 802.11 Hotspots”, *Proc. ACM MobiSys 2004*, pp. 84–97, 2004.
- [25] R. Wattenhofer, L. Li, P. Bahl, Y. Wang, “Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks”, *Proc. IEEE Infocom 2001*, pp. 1388–1397, 2001.
- [26] W. Wang, X.-Y. Li, Y. Wang, “Truthful Multicast in Selfish Wireless Networks”, *Proc. ACM MobiCom*, 2004.
- [27] J. Woodard, D. Parkes, “Strategyproof Mechanisms for Ad Hoc Network Formation”, *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.
- [28] M. Xiao, N.B. Schroff, E.K.P. Chong, “A Utility-Based Power Control Scheme in Wireless Cellular Systems”, *IEEE/ACM Trans. on Networking*, vol. 11, n. 10, pp. 210–221, 2003.
- [29] Selwyn Yuen; Baochun Li, “Strategyproof mechanisms for dynamic multicast tree formation in overlay networks”, *Wireless Networks*, Proc. IEEE Infocom 2005.
- [30] F. Xue, P.R. Kumar, “The Number of Neighbors Needed for Connectivity of Wireless Networks”, *Wireless Networks*, Vol 10, n. 2, pp. 169–181, 2004.
- [31] S. Zhong, Yang Richard Yang, J. Chen, “Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-hoc Networks”, *Proc. IEEE Infocom 2003*, pp. 1987–1997, 2003.
- [32] S. Zhong, L. Li, Y. Liu, Y.R. Yang, “On Designing Incentive-Compatible Routing and Forwarding Protocols in Wireless Ad-Hoc Networks—An Integrated Approach Using Game Theoretical and Cryptographic Techniques”, *Proc. ACM Mobicom*, Cologne, Germany, Aug 28 - Sept 2, 2005.

8. APPENDIX

The proofs are reported for the sake of completeness, and will be omitted in the final version of the paper, if accepted.

Proof of Th. 1. First, we observe that it is in the interest of a selfish node to join the network topology (otherwise it gets utility $-M$). Given this fact, how can a node maximize its utility function?

Observe that the last term in the definition of utility function (4) can be rewritten as follows:

$$pr(u, T) = \sum_{e=(u,v) \in T} pr(e).$$

The pricing rules ensures that, for any given edge e in the final topology T , both endpoints of e receive a premium at least equal to the cost of the edge. In other words, for each edge $e = (u, v) \in T$, we have $pr(e) - \frac{w_e}{2} \geq 0$.¹² Since a node in general does not know the amount of the premium for an edge it is involved in (we recall that the amount of the premium is determined by the cost of the ‘second best edge’), but it does know that it can only (although not necessarily) increase its utility by having an edge incident into it in the final topology, it follows that it is in the best interest of any node to include as many edges incidents into it as possible in the final network topology.

A node participating in IC-MST can cheat in several different ways. We show that none of this cheating possibilities leads to an increase of the node utility.

At the beginning of the protocol execution, a node (say, u) is requested to send a hello message at maximum power, including in the message the transmit power used. Node u might not send the help message at all, but in this case it would not be included in the network topology, driving its utility down to $-M$. If node u sends the message at a power less than the maximum, it exposes itself to the risk of not reaching nodes that it would have otherwise reached if sending the message at maximum power, possibly failing to connect to the network even if this would be possible. On the other hand, the weight of an edge is computed based on the received power at the receiver end, and on the transmit power as included in the hello message by the sender (not the actual transmit power). Sending the hello message with a decreased power would result in a lower received power at the receiver end, driving up the cost of all the edges incident into u . In turn, this lowers the likelihood of having many of u ’s edges in the MST, driving down its utility function. If u overdeclares the transmit power in its hello message it increases the weight of all the edges incident into u , with a negative effect on its utility. On the other hand, if u underdeclares the transmit power, all its incoming links would be compromised, as its neighbor nodes would use an incorrect (too low) transmit power to send the messages to u . Hence, u would be disconnected from the network, driving its utility down to $-M$. We also observe that u has no interest in both anticipating nor delaying the transmission of the help message, as the moment at which it joins the network topology has no effect on the payments due to the initiator, nor on the premiums received. Finally, u cannot send the hello message more than once, as, in case it would send several hello messages, its neighbors would immediately identify node u as a cheater, excluding it from the network topology (we are assuming no collusion). Thus, we have proven that a selfish node has no interests in deviating from IC-MST’s specifications during the hello message exchange phase.

Let us now consider the subsequent stages of IC-MST’s execution. Let us consider the moment in which a certain node u is first inserted into the network topology. The fact that it is inserted in the network topology depends on the weights of the edges incident into it, which are computed during the hello message exchange phase. Since we have shown above that this phase is performed correctly by a selfish node, these weights are computed correctly. Note that in principle node u might increase its utility by overdeclaring *only* the cost of the second best edge connecting it to a node in the current topology, but this is not possible, because the mechanism that we use to compute the edge weights ensures that either all the weights of the edges incident into a node are overdeclared, or all of them are underdeclared, or all of them are declared correctly. After node u is included in the network topology for the first time, it is requested to send an encrypted message to the initiator, declaring its neighbor list and corresponding edge costs. If u does not send this message, the initiator can identify u as a cheater, excluding it from the network topology. Furthermore, sending bogus information exposes node u to the risk of being identified as a cheater, as the initiator keeps track of the neighbor lists obtained from the other nodes and can easily perform a cross check (here, the assumption of no collusion between nodes is necessary). So, we have proven that it is in node u ’s best interest to behave according to IC-MST’s specifications when it is first included in the network topology.

The final case to consider is when node u is already part of the network topology, and a new node joins the network. In principle, node u might increase its utility by not reporting the cost of the second best edge, say e , incident into it (we

¹²Here, we are not constraining different edges in the max-power graph to have different weights. I.e., IC-MST is incentive compatible in the general scenario in which there exist different edges with same weight in the maxpower graph.

recall that reporting a bogus cost would expose u to the risk of being identified as a cheater and excluded from the network topology). However, also in this case the initiator will eventually receive from the node at the other end of edge e the list of neighbors, including node u and the corresponding edge cost. By cross checking this information, the initiator can easily identify u as a cheater and excluding it from the network (this is true under the assumption of no collusion). This proves that a selfish node has no interest in cheating also when it is already part of the computed topology.

To prove the theorem, it is sufficient to observe that in the last round the premiums are delivered in a secure way, as well as the payments due by the nodes in the various rounds of IC-MST's execution.

Proof of Th. 2. We recall that the social utility function is defined as follows:

$$Soc(T) = \sum_{u \in N} U(u, T) .$$

Assuming that the generated topology T is connected (otherwise at least one of the nodes' utility would be $-M$, driving down the social utility), the social utility function can be rewritten as follows:

$$Soc(T) = \sum_{u \in N} \left(- \sum_{e=(u,v) \in E_T} \frac{w_e}{2} - pay(u, T) + pr(u, T) \right) .$$

Observe that, given our definition of the pricing scheme, the sum of the payments over all the network nodes equal the sum of the premiums paid to the nodes, i.e.

$$\sum_{u \in N} (-pay(u, T) + pr(u, T)) = 0 .$$

It follows that

$$Soc(T) = \sum_{u \in N} \left(- \sum_{e=(u,v) \in E_T} \frac{w_e}{2} \right) = - \sum_{e \in E_T} w_e ,$$

implying that the function $Soc(T)$ is maximized (i.e., the social optimum is achieved) when T is the MST.

Proof of Th. 3. All the edge weights can be computed after all the network nodes have sent their hello message, i.e. with n messages overall.

During round i , a node in T_{i-1} , say u , is requested to report the weights of its best edges to nodes outside the current network topology, which can be combined into one message which is passed to the parent of u in the current topology T_{i-1} (we recall that all the intermediate topologies built by IC-MST are trees). The parent of u waits for all its children to report their weights, then it computes the best two weights out the ones received by its children and its own, and forward a unique message up in the tree. This process is repeated until all the information has been conveyed to the initiator node. Hence, at round i at most i messages are exchanged to propagate the information to the initiator, with a total of $O(n^2)$ messages during the $n-1$ rounds needed to build the network topology. It is easy to see that a similar approach can be used to collect the payments from the nodes in T_i , sending $O(n^2)$ messages overall during the $n-1$ rounds.

In the final round, the premiums are delivered to the nodes along the tree in a top-down fashion, sending $O(n)$ messages overall. It follows that the message complexity of IC-MST is $O(n^2)$.