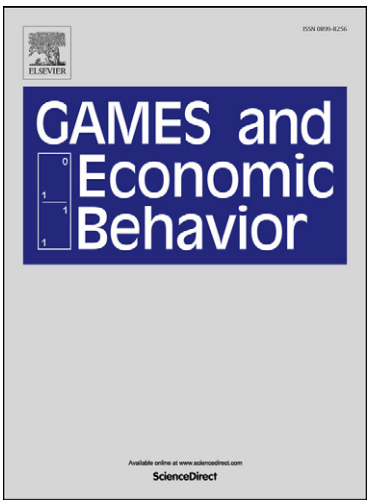

Secure message transmission on directed networks

Jérôme Renault, Ludovic Renou, Tristan Tomala

PII: S0899-8256(14)00017-7
DOI: 10.1016/j.geb.2014.01.012
Reference: YGAME 2270

To appear in: *Games and Economic Behavior*

Received date: 18 April 2012

Please cite this article in press as: Renault, J., et al. Secure message transmission on directed networks. *Games Econ. Behav.* (2014), http://dx.doi.org/10.1016/j.geb.2014.01.012

Secure message transmission on directed networks

Jérôme Renault* & Ludovic Renou† & Tristan Tomala‡

January 30, 2014

*TSE (GREMAQ, Université Toulouse 1 Capitole), 21 allée de Brienne, 31000 Toulouse, France. jerome.renault@tse-fr.eu

†Department of Economics, University of Essex, CO4 3SQ, United Kingdom. lrenou@essex.ac.uk

‡HEC Paris (Department of Economics and Decision Sciences, GREGHEC), 1 rue de la Libération, 78351 Jouy-en-Josas, France. tomala@hec.fr. Tristan Tomala acknowledges financial support of the HEC foundation. Jérôme Renault and Tristan Tomala acknowledge the support of the Agence Nationale de la Recherche, under grant ANR JEUDY, ANR-10-BLAN 0112.

Abstract

A sender wishes to transmit a secret to a receiver through a communication network, where some nodes are controlled by an adversary. We characterize the directed networks for which there exist ε-secret and ε-strongly secure communication protocols ($\forall \varepsilon > 0$): if all nodes are obedient the receiver learns the secret with probability at least $1 - \varepsilon$ and no information is leaked (secrecy), and this property is maintained under every strategy of the adversary (security). For secrecy, a necessary and sufficient condition is that there is a directed path from the sender to the receiver, and for each possible adversarial coalition A, there is an undirected path from the sender to the receiver that contains no node in A. For security, a necessary and sufficient condition is that for every possible adversarial coalition A, the graph obtained by removing all nodes in A still has the previous property.

Keywords: secure communication, protocols, communication.
JEL Classification Numbers: C72, D82

1 Introduction

The problem. This paper studies the problem of secure transmission of messages between a sender and a receiver on a network. The sender wishes to transmit a secret (private information) to the receiver through the network. The difficulty is that nodes might be curious, faulty, and malevolent (for short, adversarial), and that the identity of adversarial nodes is unknown to the sender and receiver. Nodes (players) in the network receive messages on in-going edges and send messages on out-going edges. A communication protocol specifies the messages nodes can receive and send at each round of communication (a game-form) *and* the messages nodes send at each round as a function of the messages received and sent in all previous rounds (a strategy profile). A communication protocol is *strongly secure* if it satisfies *(i)* secrecy: if the nodes correctly execute the protocol, the receiver correctly learns the secret of the sender and no information is leaked, and *(ii)* security/reliability: even if a group of nodes (an adversary) mis-execute the protocol, no information is leaked and the receiver learns the secret with arbitrary high probability. This paper characterizes the networks for which there exists a strongly secure communication protocol.

Motivation. Problems of secure transmission of messages are natural in computer science where a communication network represents agents sending messages (emails) along physical wires and adversarial players are hackers. Since the very beginning of electronic communication, security of messages has been a central concern, and there is a large literature on these issues (on communication networks and secure transmission of messages, see e.g., Dolev et al, 1993, or Franklin and Wright, 2000). The present paper contributes to this literature in considering the general class of directed networks; most of this literature considers undirected networks. We explain in greater details the connections of our work with this literature in a later section.

Yet, the primary motivation for our work is the study of game theoretic problems. Indeed, problems of secure transmission of messages are easily formalized as dynamic games. As a first example, consider a model *à la* Crawford and Sobel (1982) in which the sender and the receiver have common interests. If the sender could communicate directly with the receiver, the most efficient equilibrium would be for the sender to report truthfully his private information. Suppose now that the sender and the receiver

are distant nodes in the network and that some other players (nodes) have opposite interests. The possibility of achieving perfect communication in equilibrium is tightly related to the existence of a secure communication protocol: if such a protocol exists, then no player can alter communication by deviating, interpreting a deviation as adversarial behavior. The equivalence between secure communication protocols and fully revealing equilibria in cheap-talk games with incomplete information is studied in details in Renault and Tomala (2004).

As a second example, Monderer and Tennenholtz (1999) and Renou and Tomala (2012) study mechanism design problems when not all agents are able to directly and privately communicate with the principal[1]. The information agents like to transmit to the principal has to go through a network (the hierarchy). If one wants to implement any incentive-compatible social choice function, the problem is essentially a problem of secure transmission of messages. Intuitively, any agent must be able to securely and reliably transmit his private information to the principal. If not, either the principal receives an erroneous information and thus does not implement the right decision, or agents might have an incentive to mis-report their private information, if information about the private information of others is leaked to them.[2]

As another example, consider an infinitely repeated game. If players want to enforce a cooperative behavior (e.g., monopoly pricing), players may need to punish the deviating players upon detecting deviations from the cooperative behavior. Detecting and punishing a deviating player may prove particularly difficult if monitoring is imperfect and private, see Mailath and Samuelson (2006). For instance, suppose, that players are on a network and can only observe the actions of their neighbors. To enforce the cooperative behavior, players observing a deviation from a neighbor would need to reliably signal the identity of the deviator (the private information) to the other players in order for the punishment phase to start. Ben-Porath and Kahneman (1996) and Renault and Tomala (1998) study this problem and link the possibility of obtaining cooperative equilibria with the topology of the network. Tomala (2011) describes

[1]There is also a related literature on distributed mechanism design, see e.g. Shneidman and Parkes, 2004.

[2]Recall that the incentive compatibility of a social choice function means that no agent has an incentive to mis-report his private information when he expects others to report truthfully their private information *and* has no additional information about their private information than his prior information.

communication protocols for the identification and communication of the identity of a deviating player.

As yet another motivating problem, consider the implementation of correlated or communication equilibria by unmediated communication, see Forges (1990), Barany (1992), Ben-Porath (2003), Gerardi (2004), Abraham et al. (2006, 2008). This literature is essentially based on communication protocols possessing secrecy properties. An essential assumption in that literature is that players can communicate face to face (a complete network). To extend these results to more general communication networks, a careful study of the problem of secure transmission of messages seems a prerequisite.

Contributions. This paper generalizes previous results in the computer science literature on secure transmission of messages. Two seminal contributions are Dolev et al. (1993) and Franklin and Wright (2000). Both papers consider undirected networks where messages can travel both ways on edges. Moreover, the networks are *simple* in that they are made of parallel lines of edges. The study of general undirected networks was done by Renault and Tomala (2008) and others, see references therein. Our main contribution is to consider general *directed* networks, where information can flow in only one direction.[3]

Our central result is to characterize the directed networks for which strongly secure transmission of messages is possible. We start by assuming that for each node i, there exists a directed path from i to the receiver. This is without loss of generality. Clearly, the sender must have a directed path to the receiver if he is to transmit his secret. Suppose now that there is a non-empty set of nodes with no directed path to the receiver. All these nodes cannot send messages to the other nodes, i.e., the nodes with a directed path to the receiver, and are therefore irrelevant for the transmission of messages from the sender to the receiver. There is thus no loss of generality in removing such nodes from the network.

With this preparation done, as a first step, we study the problem of *secret* communication, that is, communication protocols such that, if all players execute the protocol obediently, then the receiver learns the secret, and the messages jointly obtained by any

[3]Note that Dolev et al (1993) and Franklin and Wright (2000) do not formalize their analysis in game-theoretic terms, while Renault and Tomala (2008) do.

adversarial coalition do not reveal any information about the content of the secret. Our first result, Theorem 1, states that a necessary and sufficient condition for secret communication is that for any possible adversarial coalition A, there exists an *undirected* path (i.e., a path in the associated undirected graph) from the sender to the receiver that does not intersect A. A distinctive feature of this result is to relate the possibility of secret communication with the topology of the associated undirected graph, i.e., the graph obtained from the original directed graph by replacing all directed edges with undirected edges. For the proof of this result, we construct a protocol that makes use of encoding techniques and show that our connectivity condition implies that only the receiver is able to decode the message correctly.

Our main theorem, Theorem 2, then states that strongly secure transmission of messages (the adversary can now deviate from the protocol) is possible if and only if for any adversarial coalition A, the network obtained by deleting all nodes in A and their adjacent edges satisfies the condition of Theorem 1. Precisely, in the network obtained by deleting all nodes in A, each remaining node has a directed path to the receiver, and for any possible adversarial coalition A', there exists an *undirected* path from the sender to the receiver that does not intersect A'.

A building block of our main proof is a secret sharing protocol (see, e.g., Shamir, 1979, and Beimel, 2011, for a recent survey). The sender (the dealer) breaks the secret into a number of shares (one share per neighbor) such that all shares are required for recovering the secret, and no information is leaked. In turn, each neighbor of the sender must deal shares of their "secret" (i.e., all messages they have received and keys generated) to their immediate neighbors, and so on. The core of the proof is to show that no adversarial coalition can obtain enough shares to learn the secret. Another aspect is that any adversarial coalition can mis-execute the protocol. From a computer science viewpoint, this corresponds to potentially dishonest dealers and shareholders and, thus, corresponds to verifiable secret sharing schemes (see, e.g., Chor et al., 1985).

There is also a large literature in computer science that studies *rational* secret sharing in game-theoretic terms (see among others Abraham et al., 2006, Gordon and Katz, 2006, Halpern and Teague, 2004, Kol and Naor, 2008). This literature provides an important bridge between computer science and game theory, since nodes can be rational entities (firms, consumers, voters, etc) that respond to incentives, as the above economic examples attest. Our work also adopts this view. We stress, however, that

the main contribution of the paper is about the topology of the network for which secure information transmission is possible, and not about the protocols for secret sharing.

An important technical result for our analysis is the *cut* lemma. Suppose that the nodes of the graph of communication are partitioned into three subsets S, A, and R such that the sender is in S, the receiver is in R, and all paths from S to R intercept A, so that A is a cut of the graph. The cut lemma states that if the histories of messages *received and sent* by A are independent of the histories of messages *received and sent* by S, so are the histories of R. A direct consequence is that if A does not learn the secret of S, then R does not either. This result is often referred to in computer science, and the usual strategy of proof is based on entropy (see e.g, Maurer, 1999). However, this type of proof does not extend easily to directed networks (see Lemma 2 below and the example following it). To contrast with, our proof uses an auxiliary three-player repeated game with an imperfect information structure that replicates the structure of observation implied by the network. This demonstrates the power of the game-theoretic language.

Related literature. This paper contributes to the literature on secure transmission of messages. Dolev et al. (1993) consider undirected networks composed of n vertex-disjoint paths from the sender to the receiver, unicast communication, and assume that an adversary controls at most k nodes.[4] These authors show that in 1-way problems, i.e., if the information flows only from the sender to the receiver, a sufficient and necessary condition for the secure transmission of information is the $3k+1$-connectivity of the network, while in 2-way problems, i.e., when the sender and the receiver "converse," a sufficient and necessary condition is the $2k+1$-connectivity of the network.[5] Similarly, considering undirected but broadcast networks, Franklin and Wright (2000) show that a necessary and sufficient condition for the secure transmission of messages is the $2k+1$-connectivity of the network. Renault and Tomala (2008) generalize Franklin and Wright's results to general undirected networks. Considering directed networks

[4]Dolev et al. distinguish between listening adversaries and disrupting adversaries. Here, we assume that adversaries are both listening and disrupting adversaries, the containment assumption.

[5]Dolev et al. actually prove stronger results: the $2k+1$-connectivity of the network is necessary and sufficient for perfect security, i.e., when with probability one the receiver correctly learns the secret and no information is leaked.

(but still composed of vertex-disjoint paths), Desmedt and Wang (2002) show that if there are $3k+1-u \geq 2k+1$ disjoint paths from the sender to the receiver and u disjoint paths from the receiver to the sender (these u paths are also disjoint from the $3k+1-u$ paths from the sender to the receiver), then perfectly secure transmission of messages can be achieved.

Our main contribution to this literature is to consider general directed networks and general adversaries. The novelty of our results is to obtain characterizations in terms of the connectivity of the *undirected graph* associated with the directed network.A closely related study is Jain (2002) who also considers general directed networks. Jain (2002) studies a variant of secret transmission —in that he assumes that nodes are obedient and that the adversary may only eavesdrop on some edges—and constructs a protocol which is similar to the one we use for proving Theorem 1. However, Jain (2002) does not consider security, i.e. the possibility that the adversary deviates from the protocol.

The paper is organized as follows. Section 2 presents formally the model and the main definitions. We give in Section 3 the characterization of secret communication, and the main result about strongly secure communication is in Section 4. Section 5 concludes the paper with comments and open problems. Some technical proofs are relegated to the Appendix.

2 Communication on networks

This section presents our model of communication on networks and defines the concepts of information-theoretic secrecy and security. We define communication protocols and strategies using concepts and terminologies borrowed from game theory.

2.1 The communication network.

A sender S and a receiver R are two distant nodes in a *directed* graph, or digraph, $\vec{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a finite set of vertices ($\{S, R\} \subseteq \mathcal{V}$) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is a finite set of edges. Each vertex is a player and directed edges represent direct communication links. For each $i \in \mathcal{V}$, we denote $D(i) = \{j : (j,i) \in \mathcal{E}\}$ the set of predecessors of i and $C(i) = \{j : (i,j) \in \mathcal{E}\}$ the set of successors. The sender privately knows a secret θ,

a realization of the random variable $\tilde{\theta}$, drawn from a finite set Θ according to a given probability distribution P (we assume $P(\tilde{\theta} = \theta) > 0$ for each $\theta \in \Theta$).

The digraph $\vec{G}$ represents the communication possibilities, i.e., a player receives messages sent by his predecessors and send messages to his successors. Communication is point-to-point, that is, a message sent by player i to player j on the edge (i, j) is private and secure: no other player can eavesdrop on the edge (i, j) or control the flow of information.

2.2 Strategies and protocols.

A communication protocol on a digraph $\vec{G}$ is informally described as follows. There are multiple rounds and, at each round, a set of players is active and send messages to their successors. Communication is assumed to be *synchronous*: at a given round, all messages are sent simultaneously, and each player is only informed of the messages he has sent and received before that round; *unicast*: a player can send different messages to different successors; and *randomized*: messages may depend on random inputs privately chosen. A protocol specifies message spaces, the strategy used by each player to generate the messages he sends, given the messages he received and his random inputs, and how the receiver decodes the secret from his messages.

More formally, players communicate for $T < \infty$ rounds. At each round $t \leq T$, player i can send a message $m_{ij}^t \in M_{ij}^t$ to player $j \in C(i)$. We assume that all message spaces are finite and that there exists a null (silent) message $m_0 \in M_{ij}^t$ for all i, for all $j \in C(i)$, for all t. A period-t history h_i^t for player i is the list of messages received and sent before round t and is an element of $H_i^t := (\times_{j \in C(i)} M_{ij}^1) \times (\times_{j \in D(i)} M_{ji}^1) \times \cdots \times (\times_{j \in C(i)} M_{ij}^{t-1}) \times (\times_{j \in D(i)} M_{ji}^{t-1})$, denote $H_i^1 := \{\emptyset\}$ the initial empty history. A period-t strategy σ_i^t for player $i \in \mathcal{V} \setminus \{S\}$ is a map from H_i^t to $\Delta(\times_{j \in C(i)} M_{ij})$.[6] A period-$t$ strategy for the sender is a map σ_S^t from $\Theta \times H_i^t$ to $\Delta(\times_{j \in C(i)} M_{ij}^t)$. A strategy σ_i for a player is a collection $(\sigma_i^1, \ldots, \sigma_i^T)$ of strategies for each period. A strategy thus defines the messages player i sends to his successors as a function of the messages he has received and sent and possibly some coin tosses. In the sequel, for any strategy profile σ, for any subset of nodes A, we write σ_A for $(\sigma_i)_{i \in A}$ and σ_{-A} for $(\sigma_i)_{i \notin A}$. Finally, $\theta_d : H_R^{T+1} \to \Theta \cup \{\text{Pb}\}$ defines the decoding function of the receiver. As a function of

[6] For any finite set X, $\Delta(X)$ denotes the set of probability distributions over X.

all the messages he has received and sent, the receiver can either output a secret or declares that there is a problem. For instance, if the receiver is confronted with two incompatible messages like "the secret is θ" and "the secret is θ^*", he might simply declare that there is a problem instead of choosing a particular secret.

Definition 1 *A communication protocol consists of a strategy profile σ and a decoding function θ_d.*

Let $\langle \sigma, \theta_d \rangle$ be a communication protocol, denote $\mathbb{P}_\sigma$ the probability distribution over histories and secrets induced by the strategy profile σ. By convention, we assume that if the history h^t has probability zero under σ, then $\mathbb{P}_\sigma(\tilde{\theta} = \theta | h^t)$ is defined as $P(\tilde{\theta} = \theta)$. Note that a profile of strategies σ together with a decoding function θ_d induce a random variable $\hat{\theta}_d$ with values on $\Theta \cup \{\text{Pb}\}$, with $\mathbb{P}_\sigma(\hat{\theta}_d = \theta) = \sum_{h_R^{T+1}:\theta_d(h_R^{T+1})=\theta} \mathbb{P}_\sigma(h_R^{T+1})$.

2.3 Adversaries.

There is a collection $\mathcal{A}$ of potential adversaries. Each $A \in \mathcal{A}$ is a subset of $\mathcal{V} \setminus \{S, R\}$. For instance, given an integer k, $\mathcal{A}$ may be the collection of all subsets of $\mathcal{V} \setminus \{S, R\}$ with at most k elements. We allow an adversary to correlate its play, i.e., we allow for strategies of the form $\tau_A^t : H_A^t \to \Delta(\times_{i \in A} \times_{j \in C(i)} M_{ij}^t)$, where $H_A^t = \cup_{i \in A} H_i^t$. An adversary A can condition its play at round t on all messages sent and received by all nodes in A up to round t. Implicitly, players in A have access to an underlying communication network to share their information and are not constrained by the network $\vec{G}$.

The goal of the sender is to send a message to the receiver without the adversary being able to learn or to manipulate the content of the message (the secret).[7]

2.4 Secrecy and security.

Definition 2 *A protocol $\langle \sigma, \theta_d \rangle$ is ε-secret if it satisfies the following requirements:*

[7]Here, we have in mind that the sender and the receiver are better off when the receiver correctly outputs the secret of the sender. For instance, suppose that upon decoding the secret to be θ', the receiver takes decision $f(\theta')$. Assume that the sender and the receiver have the utility function u with $u(f(\theta), \theta) > u(f(\theta'), \theta)$ for all (θ, θ'), while any other node has the utility function $-u$. Clearly, the receiver and the sender are better off when the receiver correctly outputs the secret, while all other nodes are worse off. If they can, they will disrupt the communication.

1. *The receiver learns the secret with probability at least* $1-\varepsilon$*, that is,* $\forall\theta\in\Theta$*,* $\mathbb{P}_\sigma(\hat{\theta}_d=\tilde{\theta}|\tilde{\theta}=\theta)\geq 1-\varepsilon$*.*

2. *No adversary gets information about the secret, that is,* $\forall A\in\mathcal{A}$*,* $\forall h_A^{T+1}$*,* $\forall\theta\in\Theta$*,* $\mathbb{P}_\sigma(\tilde{\theta}=\theta|h_A^{T+1})=P(\tilde{\theta}=\theta)$*.*

A communication protocol is ε-secret if the receiver learns the private information (secret) of the sender with arbitrarily high probability and no adversary A with unbounded computational power controlling all nodes in A (and knowing the protocol) gains additional information about the secret. Notice that the probabilities are evaluated under σ, i.e., whenever the players correctly execute the protocol. This type of adversary is referred to as "honest but curious" in the computer science literature.

We now consider a stronger requirement: the protocol must be secret whenever the players abide by the protocol, and if an adversary deviates from the protocol, the deviation is detected with arbitrarily high probability. An important motivation for this stronger requirement comes from mechanism design models (see Renou and Tomala, 2012) where the receiver is a decision maker who takes an action that affects the utility of all players, even adversarial ones. For instance, suppose that upon detecting a deviation, the receiver can take an action that imposes a large punishment on all the players. The threat of a punishment upon detection, which would happen with a large probability if the protocol has the aforementioned properties, would then deter players from deviating, i.e., to become adversarial.

Definition 3 *A protocol* $\langle\sigma,\theta_d\rangle$ *is* ε*-secret with* δ*-detection if it is* ε*-secret and satisfies the following additional requirement:*

3. *If an adversary deviates from the protocol, the receiver either correctly learns the secret or detects a problem, i.e., for all* $A\in\mathcal{A}$*, for all* τ_A*,* $\mathbb{P}_{\tau_A,\sigma_{-A}}(\hat{\theta}_d\in\{\tilde{\theta},Pb\})\geq 1-\delta$*.*

Note that if a communication protocol is secret with detection, the receiver detects with arbitrary high probability a deviation from the protocol, but may not learn the secret upon detection of the deviation. The next concept we introduce imposes that even if an adversary deviates, the protocol remains secret, i.e., the receiver still learns

the secret with high probability and no other players get additional information about the secret.

Definition 4 *A protocol $\langle \sigma, \theta_d \rangle$ is ε-strongly secure if for any adversary A and any deviation τ_A from the protocol,*

1. *The receiver learns the secret with probability at least $1-\varepsilon$, that is, $\forall \theta \in \Theta$,* $\mathbb{P}_{\tau_A, \sigma_{-A}}(\hat{\theta}_d = \tilde{\theta} | \tilde{\theta} = \theta) \geq 1 - \varepsilon$.
2. *No adversary gets information about the secret, that is, $\forall A' \in \mathcal{A}$, $\forall h_{A'}^{T+1}$, $\forall \theta \in \Theta$,* $\mathbb{P}_{\tau_A, \sigma_{-A}}(\tilde{\theta} = \theta | h_{A'}^{T+1}) = P(\tilde{\theta} = \theta)$.

In other words, for any adversary $A \in \mathcal{A}$, for any τ_A, the protocol $\langle (\tau_A, \sigma_{-A}), \theta_d \rangle$ is ε-secret.

A communication protocol is ε-strongly secure if it is ε-secret, and even if an adversary deviates from the protocol, the receiver still learns the secret with high probability and no *other* adversary gets additional information about the secret. This is stronger than the classical definition of security in computer science. Indeed, the classical definition requires that the receiver correctly learns the secret even if an adversary deviates from the protocol, that no adversary gains information by deviating, but does not require further secrecy requirement following a deviation. In particular, it might be that for the receiver to learn the secret after a deviation, other nodes have to learn it too, see the concluding example. Note that if a protocol is ε-strongly secure, then it is clearly ε-secret with ε-detection (i.e., Definition 4 implies Definition 3).

2.5 Connectivity.

Security and secrecy clearly cannot be achieved for all graphs and all adversaries. For instance, this is impossible if the adversary controls all nodes in the network and if the receiver is not a successor of the sender. We introduce now some connectivity conditions.

A *directed path* (dipath) $\vec{\gamma}$ is a finite sequence of vertices $(i_1, \ldots, i_n)$ such that $i_{k+1} \in C(i_k)$ for each $k < n$. The vertex i_1 is the origin of the path and the vertex i_n is its end-point. The digraph is strongly 1-connected from i to R, if there is a dipath with

origin i and end-point R. The digraph is strongly 1-connected to R if it is strongly 1-connected from i to R, for all $i \in \mathcal{V} \setminus \{R\}$.

We assume throughout that the digraph $\vec{G}$ is strongly 1-connected to R. Clearly, the sender must have a directed path to the receiver to be able to transmit his information. Moreover, if player $i \in \mathcal{V}\setminus\{S, R\}$ has no directed path to the receiver, then i is irrelevant for the transmission of information: either i is a sink or all directed paths starting at i terminate at a sink $i' \neq i$ or loop back to i. Strong 1-connectedness is thus without loss of generality.

With the digraph $\vec{G}$, we associate the undirected graph G obtained from $\vec{G}$ by disregarding the orientation of the edges. An undirected path (simply a path) is a finite sequence of vertices $i_1, \ldots, i_n$ such that $i_{k+1} \in C(i_k) \cup D(i_k)$ for each $k < n$.

Definition 5 *Given a collection $\mathcal{A}$ of subsets of nodes, the digraph $\vec{G}$ is* weakly $\mathcal{A}$-connected from i to R, *if for each $A \in \mathcal{A}$, there exists an undirected path from i to R with no vertex in A.*

When $\mathcal{A}$ is the collection of all subsets with at most k nodes, the undirected graph G is said to be k-connected and the digraph $\vec{G}$ is said to be weakly k-connected (see e.g. Bang-Jensen and Gutin, 2007, Chapter 1, pages 16–22).

2.6 Acyclic graphs.

The digraph is *acyclic* if each vertex appears at most once in each dipath. An important implication of acyclicity is the existence of a timing structure (i.e., an acyclic ordering).

Lemma 1 *Let $\vec{G}$ be acyclic. There exists an integer T and a function $t : \mathcal{V} \to \{1, \ldots, T\}$ such that for each $i \in \mathcal{V}$, $t(i) = 1 + \max\{t(j) : j \in D(i)\}$.*

With an acyclic graph, players do not receive feedback about the execution of the protocol. Thus, using the timing structure, we can restrict attention to protocols where each player i sends messages only at round $t(i)$, after having received messages from all his predecessors. Acyclic graphs are central to the proof of Theorem 1.

3 Secrecy

Theorem 1 *Let $\vec{G}$ be strongly 1-connected to R. The following statements are equivalent:*

1. *The graph $\vec{G}$ is weakly $\mathcal{A}$-connected from S to R.*
2. *For each $\varepsilon > 0$, there exists an ε-secret protocol.*
3. *For each $\delta > 0$, there exists a 0-secret protocol with δ-detection.*

This result gives a simple characterization of secrecy. Remarkably, only the connectivity of the associated undirected network matters (given that the networks is strongly 1-connected to R). Also, requiring detection on top of secrecy does not affect the connectivity requirements.

We prove first (2) $\Rightarrow$ (1), thereby showing the necessity of weak $\mathcal{A}$-connectedness. Then, we show that this condition is sufficient, i.e., (1) $\Rightarrow$ (3). Since (3) obviously implies (2), the proof will be complete.

3.1 Proof of necessity$[(2) \Longrightarrow (1)]$

The main tool is the fundamental next lemma. Suppose that there is a set of vertices A such that all (directed and undirected) paths from S to R intersect A, i.e. A is a *cut* of the graph. Lemma 2 states that if the histories (of messages sent and received) of the sender, are independent of the histories (of messages sent and received) of A, then the histories of the sender are also independent of the histories (of messages sent and received) of the receiver. The general statement is the following.

Lemma 2 (Cut) *Let S_1, S_2 and S_3 be three disjoint subsets of nodes such all paths from S_1 to S_3 intersect S_2, i.e., S_2 is a cut of the graph. For any σ, for any t, for any h^t, we have $\mathbb{P}_\sigma(h^t)\mathbb{P}_\sigma(h_2^t) = \mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(h_2^t, h_3^t)$, where $h_1^t = (h_i^t)_{i\in S_1}$, $h_2^t = (h_i^t)_{i\in S_2}$, and $h_3^t = (h_i^t)_{i\in S_3}$.*

The proof is in Appendix. A direct implication of the lemma, with S_1 being the sender, $S_2 = A$ and S_3 the receiver, is that the histories of messages sent and received by the receiver are independent of the secret when the histories of messages sent and received by the cut are independent.

While the intuition of the lemma is clear, we could not find a proof in the literature. In particular, let us stress that conditioning on messages *both sent and received* is crucial. For instance, it can be that the histories of messages *sent* by the cut are independent of the secret and yet the histories of messages sent and received by the receiver are not. For an example, consider Figure 1, where $\theta \in \{0,1\}$, X is a random uniform draw from $\{0,1\}$ and addition is modulo 2. The message sent by A is independent of θ (see Lemma 3) and yet the messages received by R are not (there are indeed perfectly correlated with θ).

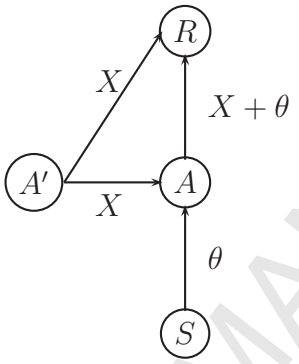

Figure 1: A is a cut.

Now, we prove [(2) $\implies$ (1)]. Let us assume that the digraph $\vec{G}$ is not weakly $\mathcal{A}$-connected from S to R, that is, there exists a set $A \in \mathcal{A}$ such that all (directed and undirected) paths from S to R intersect A, i.e. A is a cut. Let $\varepsilon < 1 - \max_{\theta' \in \Theta} P(\theta')$ and suppose by contradiction that there exists a protocol $\langle \sigma, \theta_d \rangle$ that is ε-secret. In particular, this implies that $\mathbb{P}_\sigma(\tilde{\theta} = \theta | h_A^{T+1}) = P(\tilde{\theta} = \theta)$ for each h_A^{T+1} and $\mathbb{P}_\sigma(\hat{\theta}_d = \tilde{\theta} | \tilde{\theta} = \theta) > 1 - \varepsilon$, $\forall \theta \in \Theta$.

From Lemma 2, for every terminal history $h_S^{T+1}, h_A^{T+1}, h_R^{T+1}$, we have:

$$\mathbb{P}_\sigma(h_S^{T+1}, h_A^{T+1}, h_R^{T+1}) \mathbb{P}_\sigma(h_A^{T+1}) = \mathbb{P}_\sigma(h_S^{T+1}, h_A^{T+1}) \mathbb{P}_\sigma(h_A^{T+1}, h_R^{T+1}).$$

The history of the sender h_S^{T+1} contains the secret. We fix θ in Θ, h_A^{T+1} and h_R^{T+1} and sum over all histories h_S^{T+1} that are compatible with θ and h_A^{T+1}. We obtain:

$$\mathbb{P}_\sigma(\theta, h_A^{T+1}, h_R^{T+1}) \mathbb{P}_\sigma(h_A^{T+1}) = \mathbb{P}_\sigma(\theta, h_A^{T+1}) \mathbb{P}_\sigma(h_A^{T+1}, h_R^{T+1}).$$

Using the assumption $\mathbb{P}_\sigma(\theta, h_A^{T+1}) = P(\theta) \mathbb{P}_\sigma(h_A^{T+1})$, we get:

$$\mathbb{P}_\sigma(\theta, h_A^{T+1}, h_R^{T+1}) = P(\theta) \mathbb{P}_\sigma(h_A^{T+1}, h_R^{T+1}),$$

i.e., the entire profile of messages received and sent by the cut and the receiver is independent of $\tilde{\theta}$. In particular, the profile of messages received and sent by the receiver is independent of $\tilde{\theta}$. It follows that $\hat{\theta}_d$ is independent from $\tilde{\theta}$ and, therefore, $\mathbb{P}_\sigma(\hat{\theta}_d = \tilde{\theta}) = \sum_\theta \mathbb{P}_\sigma(\hat{\theta}_d = \theta)P(\theta) < \max_{\theta\in\Theta} P(\theta) < 1 - \varepsilon$.

For secrecy, we need $\mathbb{P}_\sigma(\hat{\theta}_d = \tilde{\theta}|\tilde{\theta} = \theta) > 1 - \varepsilon$ for all θ, thus $\mathbb{P}_\sigma(\hat{\theta}_d = \tilde{\theta}) > 1 - \varepsilon$, the required contradiction. □

3.2 Proof of sufficiency[(1) ⇒ (3)]

The proof is constructive and divided into several steps. First, in subsections 3.2.1 to 3.2.4, we construct a protocol for an acyclic digraph and prove that it is 0-secret. Second, in subsection 3.2.5, we modify it in order to construct a protocol which is 0-secret with δ-detection. Finally, in subsection 3.2.6, we show how to adapt the protocol to any digraph, i.e., without the acyclicity assumption.

3.2.1 The ACY protocol.

In this section, we consider acyclic digraphs and construct a protocol, called **ACY**, which is 0-secret.

Let us encode all possible values of $\theta \in \Theta$ into binary strings. That is, we assume $\theta \in \Theta \subset \mathbb{F}^n$ for some n, where $\mathbb{F}$ denotes the finite field $\{0, 1\}$ modulo 2. Remark that $x + x = 0$ for all $x \in \mathbb{F}^n$. Throughout, a player is said to *draw a key* X if he chooses an element in $\mathbb{F}^n$ at random with equi-probability, independently of all information he may have. Remember that since the graph is acyclic, there exists a well-defined timing structure such that player i is active only at time $t(i)$.

The protocol ACY

- Sender. Chooses a unique $k_S \in C(S)$ and,
  - for each $k \in C(S) \setminus \{k_S\}$, draws X_{Sk} and sends it to k,
  - sends $\theta + \sum_{l \in D(i)} m_{li} + \sum_{k \in C(i) \setminus \{k_S\}} X_{Sk}$ to player k_S.

  The sender chooses a specific successor k_S, sends independent keys to all his successors but k_S, and sends (the sum of) all his information (type, messages received and keys) to k_S.

- Player $i \in \mathcal{V} \setminus \{S, R\}$. Chooses a unique $k_i \in C(i)$ and,
  - for each $k \in C(i) \setminus \{k_i\}$, draws X_{ik} and sends it to k,
  - sends $\sum_{l \in D(i)} m_{li} + \sum_{k \in C(i) \setminus \{k_i\}} X_{ik}$ to player k_i.

  Player i chooses a specific successor k_i, sends independent keys to all his successors but k_i, and sends (the sum of) all his information (messages received and keys) to k_i.

- Receiver. Computes $\sum_{l \in D(R)} m_{lR}$.

  Receiver computes the sum of all messages received.

Since the graph is strongly 1-connected to R, each player $i \in \mathcal{V} \setminus \{R\}$ has at least one successor, i.e., $C(i) \neq \emptyset$ for each $i \in \mathcal{V} \setminus \{R\}$, and the protocol is well defined. If a player has a unique successor, he simply sends (the sum of) all his information (and does not draw keys).

3.2.2 Example.

We illustrate the protocol **ACY** with the example in figure 2:

Let us assume that $\mathcal{A}$ is the collection of all subsets of $\{1, 2, 3, 4, 5\}$ with at most two

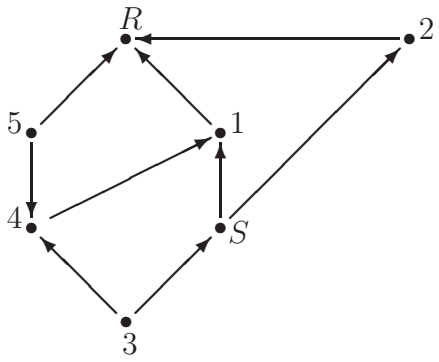

Figure 2: This graph is weakly $\mathcal{A}$-connected.

elements, so that the graph is weakly $\mathcal{A}$-connected. Notice that any directed path from the sender S to the receiver R intersects $A = \{1, 2\}$. The protocol **ACY** terminates in four rounds:

Round 1. Player 3 draws a key X_3 and sends it to the sender S and to player 4. Simultaneously, player 5 draws a key X_5 and sends it to player 4 and to the receiver R. All other players are inactive.

Round 2. Player 4 sends $X_3 + X_5$ to player 1. The sender draws a key X_1, sends the key X_1 to player 1 and sends $\theta + X_1 + X_3$ to player 2. All other players are inactive.

Round 3. Player 1 sends $X_1 + X_3 + X_5$ to the receiver R. Player 2 sends $\theta + X_1 + X_3$ to the receiver. All other players are inactive.

Round 4. The receiver computes $X_5 + (X_1 + X_3 + X_5) + (\theta + X_1 + X_3) = \theta$.

Note that each player is active at only one round and that the receiver does not need to communicate (he cannot communicate, in fact). This property of the protocol **ACY** follows from the acyclicity of the graph and the existence of the timing function (see Lemma 1).

The role of weak $\mathcal{A}$-connectivity is clear. Indeed, without the undirected path $S \leftarrow 3 \rightarrow 4 \leftarrow 5 \rightarrow R$, all paths from S to R would intersect $\{1, 2\}$ and thus the adversary would learn the secret, if the receiver is to learn it. The undirected path $S \leftarrow 3 \rightarrow 4 \leftarrow 5 \rightarrow R$ is therefore crucial. Intuitively, players 3 and 5 serve as dealers of encoding keys, which are entangled (added) by player 4. Only the receiver is then able to disentangle all keys and to decode correctly the message, as the subsequent proof shows.

3.2.3 Encoding keys and secret sharing

Before proving secrecy of this protocol, we recall a simple result about the independence of random variables (see e.g., Theorem 8.13 (p. 229) of Shoup (2008)).

Lemma 3 *Let $\mathcal{G}$ be a finite abelian group and X be a random variable uniformly distributed over $\mathcal{G}$. Let Y be a random variable in a finite set S such that X and Y are independent and let $f : S \to \mathcal{G}$ be a function. Then, the random variable $Z = X + f(Y)$ is independent from Y and uniformly distributed over $\mathcal{G}$.*

This result is well-known and the proof is straightforward. It is enough to remark that conditionally on $\{Y = y\}$, $Z = X + f(y)$ is a shift of X, and therefore is uniformly distributed for each y.

Elaborating on this result, we show that linear combinations of i.i.d. uniform random variables in $\mathbb{F}^n$ are stochastically independent if and only if they are linearly independent.

Let $(X_1, \ldots, X_K)$ be a family of i.i.d. uniform random variables in $\mathbb{F}^n$ and let $\mathcal{H}$ be the vector space (for the finite field $\mathbb{F}$) of all $\mathbb{F}^n$-valued random variables. Let also $\tilde{\theta} \in \mathcal{H}$ be (stochastically) independent from $(X_1, \ldots, X_K)$. Denote $H = \mathsf{vect}\{X_1, \ldots, X_K\}$ the sub-vector space of $\mathcal{H}$ spanned by $(X_1, \ldots, X_K)$ and $H_{\tilde{\theta}} = \mathsf{vect}\{\tilde{\theta}, X_1, \ldots, X_K\}$. Clearly, stochastic independence implies linear independence and thus, $(\tilde{\theta}, X_1, \ldots, X_K)$ are linearly independent. Conversely, we have the following:

Lemma 4

1. *Let $Y_1, \ldots, Y_L$ be L linearly independent vectors in H (i.e., linear combinations of $X_1, \ldots, X_K$). Then, $Y_1, \ldots, Y_L$ are stochastically mutually independent and uniformly distributed.*

2. *Let $Y_1, \ldots, Y_L$ be L vectors in $H_{\tilde{\theta}}$ (i.e., linear combinations of $X_1, \ldots, X_K$ and $\tilde{\theta}$) such that $\tilde{\theta} \notin \mathsf{vect}\{Y_1, \ldots, Y_L\}$. Then, $(Y_1, \ldots, Y_L)$ are jointly independent from $\tilde{\theta}$.*

Proof 1. For each $l = 1, \ldots, L$, we can write $Y_l = \sum_{k=1}^{K} a_{lk} X_k$. Since the Y_l's are linearly independent, $L \leq K$ and the rank of the matrix $A = (a_{lk})$ is L. The dimension of its kernel is thus $K - L$ and the cardinality of the kernel is $2^{n(K-L)}$. Then, denoting $X = (X_1, \ldots, X_K)$, $Y = (Y_1, \ldots, Y_K)$, we have

$$\mathbb{P}(Y = y) = \mathbb{P}(AX = y) = \mathbb{P}(AX = 0) = \frac{2^{n(K-L)}}{2^{nK}} = 2^{-nL}.$$

Therefore, Y is uniformly distributed over $(\mathbb{F}^n)^L$, as desired.

2. Assume now $\tilde{\theta} \notin \mathsf{vect}\{Y_1, \ldots, Y_L\}$. Without loss of generality, assume that the Y_l's are linearly independent (otherwise, replace by a maximal linearly independent subfamily). Since $Y_l \in H_{\tilde{\theta}}$, we write $Y_l = a_{l0}\tilde{\theta} + \sum_{k=1}^{K} a_{lk} X_k$ and set $Y_l' = Y_l - a_{l0}\tilde{\theta}$.

The Y_l''s are in H and are linearly independent. To see this, suppose to the contrary that there exists a non-trivial linear combination of $(Y_1', \ldots, Y_L')$ such that $\sum_l a_l' Y_l' = 0$. From the definition of Y_l', it follows that $\sum_l a_l' Y_l = \tilde{\theta} \sum_l a_l' a_{l0}$. If $\sum_l a_l' a_{l0} \neq 0$, this contradicts $\tilde{\theta} \notin \mathsf{vect}\{Y_1, \ldots, Y_L\}$, and if $\sum_l a_l' a_{l0} = 0$, this contradicts the assumption that the Y_l's are linearly independent.

From point 1, we conclude that the Y_l''s are stochastically independent and uniformly distributed. By construction, they are jointly independent from $\tilde{\theta}$. Then,

$$(Y_1, \ldots, Y_L) = (a_{10}, \ldots, a_{1L})\tilde{\theta} + (Y_1', \ldots, Y_L')$$

and by Lemma 3, the random vector $(Y_1, \ldots, Y_L)$ is independent of $\tilde{\theta}$ and uniformly distributed over $(\mathbb{F}^n)^L$. □

A simple method for *secret sharing* (see Shamir, 1979) is easily deduced. Informally, the aim is to "break" a secret θ into a number M of shares in such a way that one can recover the secret from all the shares but not from any subset of shares.

Lemma 5 (Secret sharing) *Let $\tilde{\theta}$ be a random variable with values in $\mathbb{F}^n$. For each integer $M > 1$, there exists a family $(X_1, \ldots, X_M)$ of random variables such that:*

1. *$\sum_{m=1}^{M} X_m = \tilde{\theta}$ almost surely.*
2. *For each m, $X_{-m} := (X_l)_{l \neq m}$ is uniformly distributed over $(\mathbb{F}^n)^{M-1}$ and independent of $\tilde{\theta}$.*

Proof Let $X_1, \ldots, X_{M-1}$ be independently and uniformly distributed over $\mathbb{F}^n$, independent of $\tilde{\theta}$, and set $X_M = \tilde{\theta} + \sum_{m=1}^{M-1} X_m$. Clearly, $\sum_{m=1}^{M} X_m = \tilde{\theta}$.

Moreover, since $X_M = X_1 + \tilde{\theta} + \sum_{m=2}^{M-1} X_m$ and $(X_1, \ldots, X_{M-1})$ are (linearly and stochastically) independent, it follows that each subfamily $X_{-m} := (X_l)_{l \neq m}$ of $M-1$ elements is linearly independent. From Lemma 4, we have that $X_{-m} := (X_l)_{l \neq m}$ are jointly independent from $\tilde{\theta}$, mutually independent and uniformly distributed. □

An important consequence is that, in the specification of the protocol **ACY**, the particular choice of a successor $k_i \in C(i)$ by player i is immaterial. From Lemma 5, all successors of a given player i are as a matter of fact treated symmetrically. That is, any *strict* subset of successors of player i, receive messages which are jointly independent from the aggregated information of player i, whereas this information can be fully recovered from the messages of *all* successors of player i. This observation allows to give the following simple and equivalent description of the protocol **ACY**.

The protocol ACY

Each players adds up his information (messages and type, if sender), breaks it into as many shares as successors, and sends one share to each successor.

3.2.4 ACY is secret

We now prove that the protocol **ACY** satisfies 0-secrecy.

Lemma 6 *Let $\vec{G}$ be acyclic, strongly 1-connected to R and weakly $\mathcal{A}$-connected from S to R. The protocol ACY on $\vec{G}$ is 0-secret.*

The proof proceeds in two claims.

Claim 1 *Under the* **ACY** *protocol, the receiver correctly learns the secret, i.e. with probability one,*

$$\sum_{l \in D(R)} m_{lR} = \tilde{\theta}$$

Proof From the definition of the protocol, $\sum_{l \in D(R)} m_{lR}$ is a linear combination of $\tilde{\theta}$ and of the keys drawn by the players. Consider a player i (possibly the sender) and a key X_{ik}, $k \neq k_i$ drawn by player i. This key is sent by i on two edges (ik and ik_i). Each other player j sends received information on exactly one edge. It follows that X_{ik} travels on exactly two directed paths with origin i. Since the graph is acyclic and strongly 1-connected, these two paths must intersect at some point (possibly at R). The player at this point adds up the messages he receives, thus X_{ik} cancels out ($X_{ik} + X_{ik} = 0$). By contrast, $\tilde{\theta}$ travels on exactly one path and is thus received by R. □

It remains to prove that no adversary $A \in \mathcal{A}$ gains additional information about θ.

Claim 2 *For any set of players A such that there exists a path from S to R disjoint from A, the messages received by players in A are jointly independent of $\tilde{\theta}$.*

It follows that if the graph is weakly $\mathcal{A}$-connected from S to R, then for each $A \in \mathcal{A}$, the messages received by players in A are jointly independent of $\tilde{\theta}$ (since then there exists an undirected path from S to R, disjoint from A). We now prove this claim.

Proof Fix a set of players A such that there exists an undirected path from S to R that contains no element of A. Denote the undirected path $\mathcal{P} := (i_0, i_1, \ldots, i_k, \ldots, i_{K+1})$, with $i_0 = S$, $i_{K+1} = R$, and $i_k \notin A$ for all $k \in \{1, \ldots, K\}$. We argue that even if the adversary A observes all the shares but the ones sent on the path $\mathcal{P}$, the adversary learns nothing about the secret.

Consider the set of players $\mathcal{P}^\circ$ on the path with no successors on the path, i.e., $\mathcal{P}^\circ := \{i_k \in \mathcal{P} : C(i_k) \cap \mathcal{P} = \emptyset\}$. Note that the path changes orientation at those players (nodes). By construction of the protocol ACY (see lemma 5), each player $i_k \in \mathcal{P}^\circ$ receives a share $X_{i_{k-1}i_k}$ from player $i_{k-1} \in \mathcal{P}$ and a share $X_{i_{k+1}i_k}$ from player $i_{k+1} \in \mathcal{P}$. (If $i_0 \in \mathcal{P}^\circ$, we let $X_{i_{-1}i_0} = \theta$.)

Assume that the adversary A observes all the shares X_{ij} such that either $i \notin \mathcal{P}$ or $j \notin \mathcal{P}$ or both. In other words, we assume that the adversary observes all possible shares but the ones sent from one player on the path $\mathcal{P}$ to another player on the path $\mathcal{P}$. It follows that the adversary observes *all* the shares of the players in $\mathcal{P}^\circ$ and, thus, learn $X_{i_{k-1}i_k} + X_{i_{k+1}i_k}$. Indeed, by construction of the protocol, player i_k computes the sum $\sum_{j \in D(i_k)} X_{j,i_k}$, breaks the sum into $|C(i_k)|$ shares and sends them to players in $C(i_k)$. Since all the shares sent to players in $C(i_k)$ as well as all the shares $(X_{ji_k})_{j \neq i_{k-1}, i_{k+1}}$ are observed by the adversary, the adversary learns $X_{i_{k-1}i_k} + X_{i_{k+1}i_k}$. For an illustration, see Figure 3.

It follows that A knows $(X_{i_{k-1}i_k} + X_{i_{k+1}i_k})_{i_k \in \mathcal{P}^\circ}$. Let i_{k^*} be the first element in $\mathcal{P}^\circ$. It is immediate to check that no linear combination of $(X_{i_{k-1}i_k} + X_{i_{k+1}i_k})_{i_k \in \mathcal{P}^\circ}$

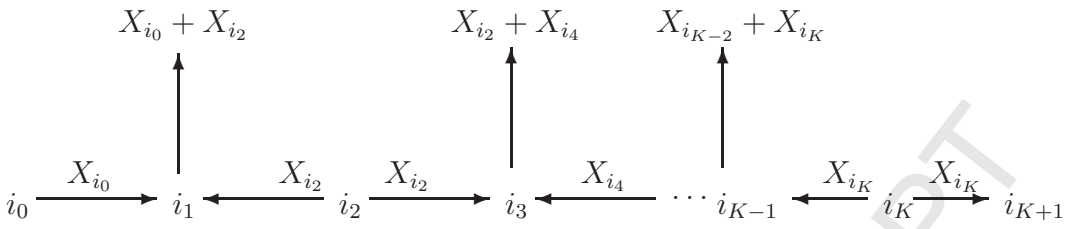

Figure 3: The undirected path $\mathcal{P}$.

gives $X_{i_{k^*-1}i_{k^*}}$.[8] Thus, the adversary cannot recover the share $X_{i_{k^*-1}i_{k^*}}$. If $i_{k^*} = i_0$, this means that the adversary learns nothing about the secret (see the second part of Lemma 5).

Alternatively, suppose that $i_{k^*} \neq i_0$. We first argue that $i_{k^*-k} \in C(i_{k^*-k-1})$ for all $k = 0, \ldots, k^*-1$, i.e., player i_{k^*-k-1} sent a share to player i_{k^*-k} for all $k = 0, \ldots, k^*-1$. In other words, the path $\mathcal{P}$ takes the form:

$$i_0 \to i_1 \to \cdots \to i_{k^*-k-1} \to i_{k^*-k} \to \cdots \to i_{k^*-1} \to i_{k^*} \leftarrow i_{k^*-1} \to \cdots \to i_{K-1} \leftarrow i_K \to i_{K+1}.$$

To see this, suppose to the contrary that there exists k such that $i_{k^*-k} \notin C(i_{k^*-k-1})$, i.e., player i_{k^*-k-1} does not send a share to player i_{k^*-k} (and, thus, receive a share from player i_{k^*-k}). Since player $i_{k^*-k-1} \notin \mathcal{P}^\circ$, this implies that player i_{k^*-k-1} must sent a share to player i_{k^*-k-2}. In turn, this implies that player i_{k^*-k-2} must send a share to player i_{k^*-k-3} (for otherwise, player i_{k^*-k-2} would be in $\mathcal{P}^\circ$). Iterating the argument, it follows that player i_1 must send a share to player i_0. Therefore, $i_0 \in \mathcal{P}^\circ$, a contradiction.

Second, we argue that the adversary does not learn the share $X_{i_{k^*-2}i_{k^*-1}}$ sent from player i_{k^*-2} to player i_{k^*-1}. Indeed, by construction of the protocol, player i_{k^*-1} computes the sum $\sum_{j \in D(i_{k^*-1})} X_{ji_{k^*-1}}$, breaks it into $|C(i_{k^*-1})|$ shares $(X_{i_{k^*-1}j})_{j \in C(i_{k^*-1})}$, and sends the share $X_{i_{k^*-1}j}$ to player $j \in C(i_{k^*-1})$. Since $i_{k^*} \in C(i_{k^*-1})$ and the adversary does not learn the share $X_{i_{k^*-1}i_{k^*}}$, it follows that the adversary learns nothing about $X_{i_{k^*-2}i_{k^*-1}}$.

Lastly, we can iterate the argument to show that the adversary learns nothing about the share $X_{i_0 i_1}$ and, consequently, about the secret of the sender. □

[8] A linear combination is the sum over a subset.

3.2.5 δ-detection.

We show now that one can build on **ACY** a 0-secret protocol with δ-detection.

Lemma 7 *Let $\vec{G}$ be acyclic, strongly 1-connected to R and weakly $\mathcal{A}$-connected, there exists a 0-secret protocol with δ-detection on $\vec{G}$.*

Proof We construct a super-protocol which consists in running in parallel a large number $N \geq 1/\delta$ of independent copies of the protocol **ACY** and where:

- Each player $i \in \mathcal{V} \setminus \{S, R\}$ must play his strategy in each copy of **ACY** and random draws are independent across copies.
- The sender selects at random one copy with probability $1/N$. Then, he inputs the secret in the selected copy and 0 in all the other copies. Without loss of generality, we assume that the set of possible secrets does not include the null binary string: $\Theta \subset \mathbb{F}^n \setminus \{(0, \ldots, 0)\}$.
- The receiver computes the output from each copy. If there is only one non-zero output, he lets the decoded secret be this single non-zero value. Otherwise, the receiver concludes that there was a deviation and declares a problem.

From the properties of the protocol **ACY**, in each copy of the protocol, any adversary $A \in \mathcal{A}$ receives messages that are independent from $\tilde{\theta}$. Moreover, these messages are mutually independent across copies. Thus, the adversary gets no information about which copy was selected by the sender. Any deviation manipulates another copy than the one selected by the sender with probability at least $1-\delta$ and is, therefore, detected since it gives at least two non-zero outputs. □

3.2.6 Dispensing with acyclicity

We now explain how to construct a secret protocol on a general digraph. The trick is to associate to the graph $\vec{G}$, an auxiliary graph $\vec{G}^{acy}$, which is strongly 1-connected, weakly $\mathcal{A}$-connected and acyclic. Then, we show how the protocol **ACY** on $\vec{G}^{acy}$ induces the desired protocol on $\vec{G}$. We start with the following observation.

Lemma 8 *There exists an acyclic and strongly 1-connected sub-graph $\vec{G}^a$ of $\vec{G}$.*

Proof. For each $i \in \mathcal{V}$, consider a shortest directed path from i to R in $\vec{G}$. Such a shortest directed path exists since $\vec{G}$ is strongly 1-connected. Let $\vec{G}^a$ be the collection of all these paths. We claim that $\vec{G}^a$ has the required properties. By construction, it is strongly 1-connected. Let us show that it is acyclic. By contradiction, assume that $\vec{G}^a$ contains the cycle $i_1 \to i_2 \to \ldots \to i_K \to i_1$. By construction, $\vec{G}^a$ is such that $C(R) = \emptyset$, i.e., there is no edge Ri for some $i \in \mathcal{V}$ in $\vec{G}^a$. It follows that the cycle does not contain the receiver. Then, there exists $k \in \{2, \ldots, K\}$ such that the shortest path from i_k to R does not follow the cycle (otherwise, R cannot be reached, a contradiction with 1-strong connectedness). Thus, the edge $i_k i_{k+1}$ is not on a shortest path from any player j to R, contradicting the construction of $\vec{G}^a$. □

Let us choose $\vec{G}^a$ to be a maximal acyclic and strongly 1-connected sub-graph of $\vec{G}$ and let $\mathcal{C} = \vec{G} \backslash \vec{G}^a$ be the set of edges of $\vec{G}$ that do not belong to $\vec{G}^a$.[9] Every edge of $\mathcal{C}$ belongs to a cycle of $\vec{G}$ and every cycle of $\vec{G}$ contains an edge in $\mathcal{C}$. Let $\vec{G}^{acy}$ be the graph obtained from $\vec{G}$ by replacing each edge ij in $\mathcal{C}$ by two edges: $i(j)i$ and $i(j)j$, where $i(j)$ is a fictitious player who is a duplicate of player i. That is, if ij is in $\mathcal{C}$:

$$i \to j \text{ is replaced by } i \leftarrow i(j) \to j.$$

The edges of $\vec{G}^a$ are unchanged. See Figure 4 for an example.

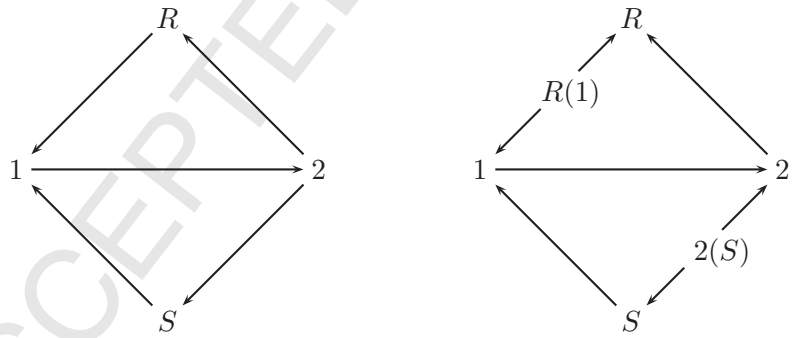

Figure 4: A cyclic graph $\vec{G}$ and the associated acyclic $\vec{G}^{acy}$

Claim 3 *$\vec{G}^{acy}$ is strongly 1-connected, weakly $\mathcal{A}$-connected from S to R and acyclic.*

[9] I.e. $\vec{G}^a$ is not a proper subgraph of a strongly 1-connected sub-graph of $\vec{G}$.

Proof. Each "regular" player i has a directed path to R in $\vec{G}^a$ by construction. Since the fictitious player $i(j)$ is directly connected to i, he also has a path to the designer by strong 1-connectedness of $\vec{G}$. Weak $\mathcal{A}$-connectedness is clearly preserved by the transformation. Let us show that $\vec{G}^{acy}$ is acyclic. Assume that $\vec{G}^{acy}$ contains a cycle. By our construction, each fictitious player has only out-going edges, thus cannot belong to a cycle. This implies that the cycle was already a cycle in $\vec{G}$ and, therefore, it should contain an edge which belongs to $\mathcal{C}$. This is a contradiction because edges in $\mathcal{C}$ no longer appear in $\vec{G}^{acy}$. □

Now, we construct a secret protocol on $\vec{G}$ by emulating the protocol **ACY** on $\vec{G}^{acy}$. The timing is the one given by the acyclic structure of $\vec{G}^{acy}$. Each player i who has duplicates in $\vec{G}^{acy}$ should play first for his duplicates (at the first round since duplicates have no predecessors) and then at round $t(i)$ prescribed by the timing structure of $\vec{G}^{acy}$. There, he should treat the messages he did choose for his duplicates as messages received by predecessors.

The secrecy of **ACY** on $\vec{G}^{acy}$ clearly implies that the induced protocol is secret on $\vec{G}$ (if a duplicated player belongs to the adversary, the messages sent by his duplicates are independent of the secret and all other messages and thus convey no information).

As in the previous section, one obtains a protocol with δ-detection by running in parallel a large number of independent copies of this protocol.

4 Security

We now present our main characterization for strong security (the adversary can now deviate from the protocol). Denote 0_A the strategy of the adversary A, which consists in sending the null message, regardless of the history.

Theorem 2 *The following statements are equivalent:*

1. *For each $A \in \mathcal{A}$, the graph $\vec{G} \setminus A$ contains a sub-graph that is strongly 1-connected to R, and weakly $\mathcal{A}$-connected from S to R.*

2. *For any $\varepsilon > 0$, there exists an ε-strongly secure protocol.*

3. *For any $\varepsilon > 0$, there exists a protocol $\langle \sigma, \theta_d \rangle$ such that $\langle (0_A, \sigma_{-A}), \theta_d \rangle$ is ε-secret for any $A \in \mathcal{A}$.*

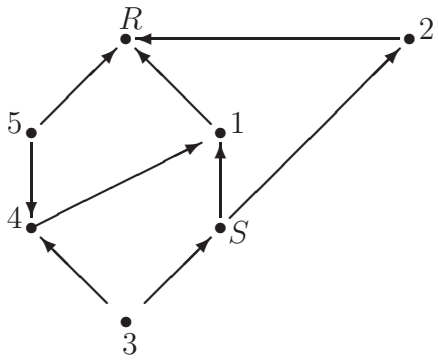

Figure 5: An example

In words, the connectivity conditions we obtain are as follows. Firstly, for any adversary A, there must exists a directed path from S to R in $\vec{G} \setminus A$. Secondly, consider the sub-graph of $\vec{G} \setminus A$ made of all players in $\vec{G} \setminus A$ who have a directed path to R. Then, this sub-graph must be weakly $\mathcal{A}$-connected, namely for any $A' \in \mathcal{A}$, there must be an undirected path in this sub-graph that does not intersect $A \cup A'$.

To get some intuition for Theorem 2, consider the simple example introduced in Figure 5.

Assume that the adversary is $A = \{1, 2\}$. It is clear that if players 1 and 2 stop communicating with the receiver, then the sender cannot communicate at all with the receiver and thus cannot communicate his secret. Thus, for any adversary A, we need at least one directed path from the sender to the receiver that does not intersect A. However, this is not sufficient to guarantee secrecy. From Theorem 1, for each adversary A', there must exist a path from the sender to the receiver that does not intersect A' to guarantee secrecy. So, if the adversary A stops communicating, we would need the sub-graph $\vec{G} \setminus A$ to be weakly $\mathcal{A}$-connected if we want achieve strong security. Theorem 2 states that it is not only a sufficient condition, but also a necessary one. Moreover, Theorem 2 also states that simple faults (i.e., to simply stop communicating) are the worst possible deviations.

We prove first sufficiency (1) $\Rightarrow$ (2), then necessity (3) $\Rightarrow$ (1), (2) $\Rightarrow$ (3) is obvious.

4.1 Sufficiency $[(1) \Longrightarrow (2)]$

Fix $\varepsilon > 0$ and set $\varepsilon' = \varepsilon/|\mathcal{A}|$. Since for each $A \in \mathcal{A}$, the graph $\vec{G} \setminus A$ contains a sub-graph that is strongly connected and weakly $\mathcal{A}$-connected from S to R, Theorem

1 applies. Therefore, there exists a protocol $\langle \sigma^A, \theta_d^A \rangle$ on the strongly connected and weakly $\mathcal{A}$-connected sub-graph of $\vec{G} \setminus A$, which guarantees 0-secrecy with ε'-detection.

Consider the protocol $\langle \sigma, \theta_d \rangle$ which consists in running in parallel all protocols $\langle \sigma^A, \theta_d^A \rangle$, $A \in \mathcal{A}$. The strategy of each $i \in \mathcal{V}$ is $\sigma_i := (\sigma_i^A)_{\{A: i \in A\}}$. The decoding rule is such that as soon as one protocol $\langle \sigma^A, \theta_d^A \rangle$ succeeds, i.e., does not end with a problem, the receiver selects the output from one of the successful protocols. Precisely, fix a linear order $<$ on $\mathcal{A}$, and for each terminal history h_R^{T+1} of the receiver, denote $A^*(h_R^{T+1})$ the least (according to $<$) $A \in \mathcal{A}$ such that $\theta_d^A(h_R^{T+1}) \in \Theta$, i.e., the first adversary for which the receiver does not detect a problem in the protocol $\langle \sigma^A, \theta_d^A \rangle$. The decoding function θ_d is then $\theta_d(h_R^{T+1}) = \theta_d^{A^*(h_R^{T+1})}(h_R^{T+1})$ for each h_R^{T+1}.

This protocol is 0-secret by construction: each sub-protocol $\langle \sigma^A, \theta_d^A \rangle$ is, and these sub-protocols are mutually independent. An adversary cannot learn information about the secret by deviating. However, he could perturb the decoding and induce the receiver to output a wrong value. Fix an adversary $A \in \mathcal{A}$. Since no node in A is active in the protocol $\langle \sigma^A, \theta_d^A \rangle$, the adversary cannot affect the outcome of $\langle \sigma^A, \theta_d^A \rangle$. It follows that the above decoding function is well-defined. However, there is no guarantee that $A = A^*(h_R^{T+1})$ for each h_R^{T+1}. Assume that the adversary manipulates $\langle \sigma^B, \theta_d^B \rangle$ with $B \neq A$, in such a way that the receiver outputs a wrong value $\theta_d^B \neq \tilde{\theta}$. The probability of this event is at most ε'. In this case, if $A^*(h_R^{T+1}) = B$, the receiver outputs a wrong value from the full protocol. Thus, the probability of wrong output is at most the probability that there exists $B \in \mathcal{A}$ such that $A^*(h_R^{T+1}) = B$ and $\theta_d^B \neq \tilde{\theta}$, which is at most $|\mathcal{A}|\varepsilon' = \varepsilon$. Thus, $\langle \sigma, \theta_d \rangle$ is ε-strongly secure.

4.2 Necessity [(3) $\Rightarrow$ (1)]

The key of the proof is the following lemma.

Lemma 9 *Consider two disjoint subsets of nodes U and V. Let $\vec{U}$ be the set of nodes $i \in \mathcal{V} \setminus U$ such that all directed paths from i to some $j \in V$, go through U.*

For all $\tau_{\vec{U}}$, $\tilde{\tau}_{\vec{U}}$, $\sigma_{-(U \cup \vec{U})}$, and for every history h_V^{T+1} of V, we have

$$\mathbb{P}_{0_U, \tau_{\vec{U}}, \sigma_{-(U \cup \vec{U})}}(h_V^{T+1}) = \mathbb{P}_{0_U, \tilde{\tau}_{\vec{U}}, \sigma_{-(U \cup \vec{U})}}(h_V^{T+1}).$$

The proof is in Appendix. The intuitive meaning of this lemma is the following. Nodes in $\vec{U}$ are separated from V by U, in the sense that messages from nodes in $\vec{U}$

to recipients in V are "filtered" by nodes in U. If U adopts the strategy 0_U, then all nodes in U send no message and therefore communication from $\vec{U}$ to V is disrupted. Consequently, the messages received by V do not depend on the strategy used by $\vec{U}$.

We prove now [(3) $\Rightarrow$ (1)] by contradiction. Suppose that there exists $A^* \in \mathcal{A}$ such that $\vec{G} \setminus A^*$ does not contain a sub-graph that is strongly connected and weakly $\mathcal{A}$-connected. Let $\vec{A^*}$ the set of all nodes $i \in \mathcal{V} \setminus A^*$ such that all *directed* paths from i to R go through A^* (i.e., in $\vec{G} \setminus A^*$, the nodes in $\vec{A^*}$ have no directed paths to the receiver R).

Assume first $S \in \vec{A^*}$, and apply Lemma 9, with $U = A^*$ and $V = \{R\}$. It follows that under the protocol $\langle (0_{A^*}, \sigma_{-A^*}), \theta_d \rangle$, the history of V, i.e., the receiver, does not depend on the value of $\tilde{\theta}$ and thus $\langle (0_{A^*}, \sigma_{-A^*}), \theta_d \rangle$ cannot be ε-secret, for any ε.

Assume now $S \notin \vec{A^*}$. By construction of $\vec{A^*}$, the sub-graph $\vec{G} \setminus (A^* \cup \vec{A^*})$ of $\vec{G} \setminus A^*$ is (maximally) strongly connected and includes the sender S. Therefore, there must exist $A \in \mathcal{A}$ such that all paths from S to R intersect A, i.e., the sub-graph $\vec{G} \setminus (A^* \cup \vec{A^*})$ is not weakly $\mathcal{A}$-connected. Note that $\vec{A^*} \cup A^* \cup A$ is cut of the graph $\vec{G}$, i.e., all paths from S to R intersect $\vec{A^*} \cup A^* \cup A$. If $\vec{A^*} \cup A^* \cup A \in \mathcal{A}$, the result follows from the necessity part of Theorem 1, so let us assume that this is not the case.

We apply Lemma 9 with $U = A^*$, $\vec{U} = \vec{A^*}$ and $V = \{R\} \cup A$. It follows that $\mathbb{P}_{0_{A^*}, \tau_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})}}(h_R^{T+1}, h_A^{T+1})$ does not depend on $\tau_{\vec{A^*}}$ and thus, we may assume without loss of generality that $\tau_{\vec{A^*}} = 0_{\vec{A^*}}$. (Note that there is no directed path from a node in $\vec{A^*}$ to a node in A that does not intersect A^*, since any node in A has a directed path to R in $G \setminus (\vec{A^*} \cup A^*)$.) We claim that there cannot exist an ε-strongly secure protocol. By contradiction, suppose that for $\varepsilon > 0$, there exists a protocol $\langle \sigma, \theta_d \rangle$ which is ε-strongly secure. This means that

$$\mathbb{P}_{0_{A^*}, \tau_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})}, \theta_d}(\tilde{\theta} = \theta | h_A^{T+1}) = P(\tilde{\theta} = \theta)$$

and

$$\forall \theta \in \Theta, \mathbb{P}_{0_{A^*}, \tau_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})}, \theta_d}(\hat{\theta}_d = \tilde{\theta} | \tilde{\theta} = \theta) > 1 - \varepsilon.$$

From Lemma 9, this implies

$$\mathbb{P}_{0_{A^*}, 0_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})}, \theta_d}(\tilde{\theta} = \theta | h_A^{T+1}) = P(\tilde{\theta} = \theta)$$

and

$$\forall \theta \in \Theta, \mathbb{P}_{0_{A^*}, 0_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})}, \theta_d}(\hat{\theta}_d = \tilde{\theta} | \tilde{\theta} = \theta) > 1 - \varepsilon.$$

This means that the protocol induced by $\langle 0_{A^*}, 0_{\vec{A^*}}, \sigma_{-(A^* \cup \vec{A^*})} \rangle$ on the restricted graph $\vec{G} \setminus (A^* \cup \vec{A^*})$ is ε-secret. This contradicts Theorem 1 as $\vec{G} \setminus (A^* \cup \vec{A^*})$ is strongly connected but not weakly $\mathcal{A}$-connected. The proof is thus complete.

5 Open questions

5.1 A weaker notion of security

As already alluded to, our notion of strong security is more stringent that the classical definition of security, which is as follows.

Definition 6 *A protocol $\langle \sigma, \theta_d \rangle$ is ε-secure if it satisfies the following requirements:*

1. *The receiver learns the secret with probability at least $1 - \varepsilon$, i.e., for all $A \in \mathcal{A}$, for all τ_A, $\mathbb{P}_{\tau_A, \sigma_{-A}}(\hat{\theta}_d = \tilde{\theta}) \geq 1 - \varepsilon$.*

2. *No adversary gets information about the secret, i.e., for all $A \in \mathcal{A}$, for all τ_A, for all h_A^{T+1}, $\mathbb{P}_{\tau_A, \sigma_{-A}}(\tilde{\theta} = \theta | h_A^{T+1}) = P(\tilde{\theta} = \theta)$.*

A protocol is ε-secure if the receiver learns the secret with arbitrarily high probability, regardless of the behavior of adversary. Moreover, no adversary can gain information by deviating. However, if an adversary deviates from the protocol, information may be leaked to *other* players. In other words, by contrast with the concept of strong security, secrecy is not maintained when adversaries deviate from the protocol.

To see the difference between security and strong security, consider the graph $\vec{G}$ below with $\mathcal{A}$ the set of all singletons. The graph $\vec{G} \setminus \{3\}$ does not contain a sub-graph that is strongly connected and weakly $\mathcal{A}$-connected, while all others do. From Theorem 2, there is therefore no strongly ε-secure protocol on $\vec{G}$ for all $\varepsilon > 0$. Yet, we claim that ε-security can be achieved.

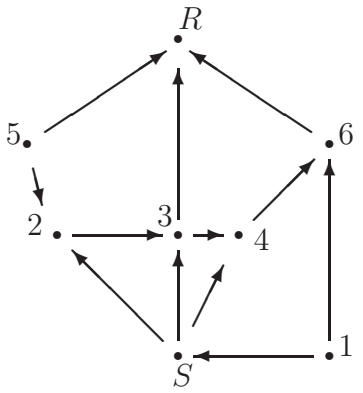

$\vec{G}$

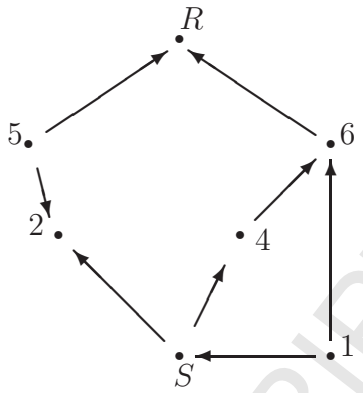

$\vec{G} \setminus \{3\}$

For any $\varepsilon > 0$, we construct a protocol, which consists of six sub-protocols run in parallel, one for each possible adversary. For each adversary $A \neq \{3\}$, we run the protocol **ACY** with ε-detection on the sub-graph $\vec{G} \setminus A$. Since for each $A \neq \{3\}$, the $\vec{G} \setminus A$ is strongly 1-connected and weakly A-connected, Theorem 1 guarantees that the protocol is 0-secret with ε-detection. For the adversary $\{3\}$, the sub-protocol is described as follows:

Round 1. Player 1 draws X_1 and sends it to the sender S and player 6. Simultaneously, player 5 draws X_5 and sends it to the receiver R and player 2. All other players are inactive.

Round 2. The sender S draws Z_S and Z'_S and sends $(\theta+X_1, Z_S, Z'_S)$ to player 4 and (Z_S, Z'_S) to player 2. All other players are inactive.

Round 3. Player 2 computes $Y_5 = Z_S X_5 + Z'_S$ and sends (X_5, Y_5) to player 3. All other players are inactive.

Round 4. Player 3 forwards (X_5, Y_5) to player 4. All other players are inactive.

Round 5. Let (X'_5, Y'_5) be the message received from player 3 by player 4. Player 4 tests whether $Y'_5 = Z_S X'_5 + Z'_S$.

- If the test succeeds, he sends $(\theta + X_1 + X'_5, \text{ok})$ to player 6.
- If the test fails, he sends $(\theta + X_1, \text{problem})$ to player 6.

All other players are inactive.

Round 6.
- If player 6 receives a message with "ok," he sends $(\theta + X_1 + X_5' + X_1, \text{ok})$ to the receiver R.
- If player 6 receives a message with "problem," he sends $(\theta + X_1 + X_1, \text{problem})$ to the receiver R.

All other players are inactive.

Round 7.
- If the receiver R receives a message with "ok," he decodes $(\theta + X_1 + X_5' + X_1) + X_5$.
- If the receiver R receives a message with "problem," he decodes $(\theta + X_1 + X_1)$.

All other players are inactive.

Clearly, if all players abide by the sub-protocol, the receiver correctly learns the secret. Moreover, if player 3 deviates from the sub-protocol, the deviation is detected whenever $Y_5' \neq Z_S X_5' + Z_S'$ (an authentication test). The probability of detection can be made arbitrarily large, in particular larger than $1 - \varepsilon$, by taking keys in $\mathbb{F}^n$ for n large enough.

Finally, the receiver fixes an arbitrary order on the six sub-protocols and decodes the secret according to the first sub-protocol for which no deviation is detected. If the receiver detects a deviation in all sub-protocols, he decodes the secret according to the sub-protocol constructed for the adversary $\{3\}$. By construction, the probability that an adversary deviates from the protocol and is not detected is at most ε, so that we indeed achieve ε-security. However, we do not achieve strong ε-security. To see this, suppose that the adversary is player 3. Note that he is active in *all* sub-protocols. If he deviates in all sub-protocols, the deviation is detected with probability at least $1 - \varepsilon$, in which case the receiver correctly learns the secret from player 6 (in the execution of the sub-protocol constructed for 3). However, in this case, player 6 also learns the secret, which violates strong ε-security.

An open problem is to characterize the directed graphs for which ε-security can be achieved for all $\varepsilon > 0$.

5.2 Perfect security

Another open problem is to characterize the networks for which 0-secrecy and 0-strong security can be achieved. To formulate a guess, let us assume that $\mathcal{A}$ is the set of

coalitions with at most k elements, and compare our conditions with those of Dolev et al. (1993).[10] We recall that they only treat networks with vertex disjoint paths, that is, graphs made of parallel lines. Yet, their conditions give intuitions for the general case.

The condition in our main theorem is similar to weak $2k+1$-connectivity: removing $2k$ nodes does not disconnect the sender from the receiver *in the undirected graph*. This is analogous to the result of Dolev et al. regarding undirected, i.e. 2-way networks, where $2k+1$-connectivity is required. For 1-way networks Dolev et al. show that $3k+1$-connectivity is required. By analogy, we conjecture the following:

Perfect security can be achieved if and only if for each pair $A, A' \in \mathcal{A}$, the graph $\vec{G} \setminus A \cup A'$ contains a sub-graph that is strongly connected and weakly $\mathcal{A}$-connected from S to R.

6 Appendix

Lemma 2 *Let S_1, S_2 and S_3 be three disjoint subsets of nodes such all paths from S_1 to S_3 intersect S_2, i.e., S_2 is a cut of the graph. For any σ, for any t, for any h^t, we have $\mathbb{P}_\sigma(h^t)\mathbb{P}_\sigma(h_2^t) = \mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(h_2^t, h_3^t)$, where $h_1^t = (h_i^t)_{i\in S_1}$, $h_2^t = (h_i^t)_{i\in S_2}$, and $h_3^t = (h_i^t)_{i\in S_3}$.*

Proof Let $\vec{G}$ be a directed graph and G the undirected graph obtained from $\vec{G}$. Assume that G is connected. Consider three disjoint subsets S_1, S_2 and S_3 of players such that all paths from players in S_1 to players in S_3 intersects S_2, i.e., S_2 is a cut of the graph G. Let S_1^* (resp., S_3^*) be the connected component of $G \setminus S_2$ that includes S_1 (resp., S_3). Since S_2 is a cut, we have that $S_1^* \cap S_3^* = \emptyset$ and if $i \notin S_1 \cup S_2 \cup S_3$, then either $i \in S_1^*$ or $i \in S_3^*$. And, of course, all paths from S_1^* to S_3^* intersect S_2.

We prove Lemma 2 with $S_1 = S_1^*$ and $S_3 = S_3^*$. The complete proof follows easily from this case by summing over the relevant histories $h^t_{S_1^* \setminus S_1}$ and $h^t_{S_3^* \setminus S_3}$.

So, we want to prove that for any σ, for any t, for any h^t, we have $\mathbb{P}_\sigma(h^t)\mathbb{P}_\sigma(h_2^t) = \mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(h_2^t, h_3^t)$, where $h_1^t = (h_i^t)_{i\in S_1}$, $h_2^t = (h_i^t)_{i\in S_2}$, and $h_3^t = (h_i^t)_{i\in S_3}$.

We first prove Lemma 2 for an auxiliary game with three players: 1, 2 and 3. We may think of player i in that auxiliary game as a representative of the players in the

[10] Dolev et al. study, among others, the problem of 0-secrecy and 0-security.

set S_i. The auxiliary game has T stages with imperfect observation.

At each stage t, player 1 chooses an action (a, α), player 2 an action (β, b, γ), and player 3 an action (c, δ). At the end of stage t, player 1 observes β, player 2 observes α and δ, and player 3 observes γ.

The action a corresponds to the messages sent by nodes in S_1 to nodes in S_1, while the action α corresponds to the messages sent by nodes in S_1 to S_2. Similarly, the action β corresponds to the messages sent by nodes in S_2 to nodes in S_1, the action b to the messages sent by nodes in S_2 to S_2 and the action γ to the messages sent by nodes in S_2 to S_3. Lastly, the action c corresponds to the messages sent by nodes in S_3 to S_3 and the action δ to the messages sent by nodes in S_3 to nodes in S_2.

A t-period history is a vector $h^t = (x_1, ..., x_t)$ with $x_s = (a_s, \alpha_s, \beta_s, b_s, \gamma_s, c_s, \delta_s)$ the actions played at stage s, $s \leq t$. Denote h_1^t player 1's history at stage t, i.e., $h_1^t = (a_s, \alpha_s, \beta_s)_{s \leq t}$. Similarly, $h_2^t = (\alpha_s, \beta_s, b_s, \gamma_s, \delta_s)_{s \leq t}$, and $h_3^t = (\gamma_s, c_s, \delta_s)_{s \leq t}$.

The proof is by induction on t. Let $t = 1$ and consider the history $h^1 = (a_1, \alpha_1, \beta_1, b_1, \gamma_1, c_1, \delta_1)$. We have

$\mathbb{P}_\sigma(h^1) = \sigma_1(a_1, \alpha_1)\sigma_2(\beta_1, b_1, \gamma_1)\sigma_3(c_1, \delta_1)$,
$\mathbb{P}_\sigma(h_2^1) = \sigma_1(\alpha_1)\sigma_2(\beta_1, b_1, \gamma_1)\sigma_3(\delta_1)$,
$\mathbb{P}_\sigma(h_1^1, h_2^1) = \sigma_1(a_1, \alpha_1)\sigma_2(\beta_1, b_1, \gamma_1)\sigma_3(\delta_1)$,
$\mathbb{P}_\sigma(h_3^1, h_2^1) = \sigma_3(c_1, \delta_1)\sigma_2(\beta_1, b_1, \gamma_1)\sigma_1(\alpha_1)$,
which establishes the statement for $t = 1$.

Let us suppose that the statement is true for t, i.e., $\mathbb{P}_\sigma(h^t)\mathbb{P}_\sigma(h_2^t) = \mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(h_2^t, h_3^t)$. Consider the history $h^{t+1} = (h^t, a_{t+1}, \alpha_{t+1}, \beta_{t+1}, b_{t+1}, \gamma_{t+1}, c_{t+1}, \delta_{t+1})$, with

$$h^t = (a_s, \alpha_s, \beta_s, b_s, \gamma_s, c_s, \delta_s)_{s \leq t}.$$

Firstly, we have

$$\begin{aligned}
\mathbb{P}_\sigma(h^{t+1}) &= \mathbb{P}_\sigma(h^t)\mathbb{P}_\sigma(a_{t+1}, \alpha_{t+1}, \beta_{t+1}, b_{t+1}, \gamma_{t+1}, c_{t+1}, \delta_{t+1}|h^t), \\
&= \mathbb{P}_\sigma(h^t)P(a_{t+1}, \alpha_{t+1}|h^t)\mathbb{P}_\sigma(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h^t)\mathbb{P}_\sigma(c_{t+1}, \delta_{t+1}|h^t), \\
&= \mathbb{P}_\sigma(h^t)\sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\sigma^3(c_{t+1}, \delta_{t+1}|h_3^t). \quad (1)
\end{aligned}$$

Secondly, we have that

$$\begin{aligned}
\mathbb{P}_\sigma(h_1^{t+1}, h_2^{t+1}) &= \sum_{c'_1,...,c'_{t+1}} \mathbb{P}_\sigma(h_1^{t+1}, h_2^{t+1}, c'_1, ..., c'_{t+1}) \\
&= \sum_{c'_1,..,c'_{t+1}} \big(\mathbb{P}_\sigma(h_1^t, h_2^t, c'_1, ..., c'_t)\sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t) \\
&\quad \sigma_3(c'_{t+1}, \delta_{t+1}|\gamma_1, .., \gamma_t, \delta_1, .., \delta_t, c'_1, .., c'_t)\big)\,, \\
&= \sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t) \\
&\quad \sum_{c'_1,...,c'_{t+1}} \mathbb{P}_\sigma(h_1^t, h_2^t, c'_1, ..., c'_t)\sigma_3(c'_{t+1}, \delta_{t+1}|h_1^t, h_2^t, c'_1, ..., c'_t), \\
&= \sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t) \\
&\quad \sum_{c'_1,...,c'_t} \mathbb{P}_\sigma(h_1^t, h_2^t, c'_1, ..., c'_t)\sigma_3(\delta_{t+1}|h_1^t, h_2^t, c'_1, ..., c'_t), \\
&= \sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\mathbb{P}_\sigma(h_1^t, h_2^t, \delta_{t+1}), \\
&= \sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_1^t, h_2^t).
\end{aligned}$$

We also have that $\mathbb{P}_\sigma(\delta_{t+1}|h_1^t, h_2^t) = \sum_{c'_1,...,c'_t} \mathbb{P}_\sigma(c'_1, ..., c'_t|h_1^t, h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_1^t, h_2^t, c'_1, ..., c'_t)$. Moreover, it follows from the induction hypothesis that $\mathbb{P}_\sigma(c'_1, ..., c'_t|h_1^t, h_2^t) = \mathbb{P}_\sigma(c'_1, ..., c'_t|h_2^t)$ and from the definition of player 3 strategy that $\mathbb{P}_\sigma(\delta_{t+1}|h_1^t, h_2^t, c'_1, ..., c'_t) = \mathbb{P}_\sigma(\delta_{t+1}|h_2^t, c'_1, ..., c'_t)$. Consequently, we have

$$\mathbb{P}_\sigma(\delta_{t+1}|h_1^t, h_2^t) = \sum_{c'_1,...,c'_t} \mathbb{P}_\sigma(c'_1, ..., c'_t|h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_2^t, c'_1, ..., c'_t) = \mathbb{P}_\sigma(\delta_{t+1}|h_2^t).$$

Finally, it follows that

$$\mathbb{P}_\sigma(h_1^{t+1}, h_2^{t+1}) = \sigma_1(a_{t+1}, \alpha_{t+1}|h_1^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_2^t). \quad (2)$$

Thirdly, exchanging the role of players 1 and 3, we obtain

$$\mathbb{P}_\sigma(h_3^{t+1}, h_2^{t+1}) = \sigma_3(c_{t+1}, \delta_{t+1}|h_3^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\mathbb{P}_\sigma(h_3^t, h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|h_2^t). \quad (3)$$

From equations (1), (2) and (3), we obtain:

$$\frac{\mathbb{P}_\sigma(h_3^{t+1}, h_2^{t+1})\mathbb{P}_\sigma(h_1^{t+1}, h_2^{t+1})}{\mathbb{P}_\sigma(h^{t+1})} = \frac{\mathbb{P}_\sigma(h_1^t, h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_2^t)\mathbb{P}_\sigma(h_3^t, h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|h_2^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)}{\mathbb{P}_\sigma(h^t)}.$$

From the induction hypothesis, we have:

$$\frac{\mathbb{P}_\sigma(h_2^{t+1}, h_3^{t+1})\mathbb{P}_\sigma(h_1^{t+1}, h_2^{t+1})}{\mathbb{P}_\sigma(h^{t+1})} = \mathbb{P}_\sigma(h_2^t)\sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|h_2^t). \tag{4}$$

It remains to show that the right-hand side of Equation 4 is actually $\mathbb{P}_\sigma(h_2^{t+1})$.

To save space, denote $\Sigma_2 = \sigma_2(\beta_{t+1}, b_{t+1}, \gamma_{t+1}|h_2^t)$ and let us compute $\mathbb{P}_\sigma(h_2^{t+1}|h_2^t)$.

$$\mathbb{P}_\sigma(h_2^{t+1}|h_2^t) =$$
$$\sum_{a'_1,..,a'_t,c'_1,...,c'_t} \mathbb{P}_\sigma(a'_1, .., a'_t, c'_1, ..., c'_t|h_2^t)\mathbb{P}_\sigma(h_2^{t+1}|a'_1, .., a'_t, c'_1, .., c'_t, h_2^t) =$$
$$\sum_{a'_1,..,a'_t,c'_1,...,c'_t} \mathbb{P}_\sigma(a'_1, .., a'_t, c'_1, ..., c'_t|h_2^t)\Sigma_2\mathbb{P}_\sigma(\delta_{t+1}|a'_1, .., a'_t, c'_1, ..., c'_t, h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|a'_1, .., a'_t, c'_1, .., c'_t, h_2^t).$$

Finally, it follows from the induction hypothesis that

$$\mathbb{P}_\sigma(a'_1, .., a'_t, c'_1, ..., c'_t|h_2^t) = \mathbb{P}_\sigma(a'_1, .., a'_t|h_2^t)\mathbb{P}_\sigma(c'_1, ..., c'_t|h_2^t),$$

and from the definition of player 2 strategies that

$$\mathbb{P}_\sigma(\delta_{t+1}|a'_1, .., a'_t, c'_1, .., c'_t, h_2^t) = \mathbb{P}_\sigma(\delta_{t+1}|c'_1, .., c'_t, h_2^t),$$

and

$$\mathbb{P}_\sigma(\alpha_{t+1}|a'_1, .., a'_t, c'_1, .., c'_t, h_2^t) = \mathbb{P}_\sigma(\alpha_{t+1}|a'_1, .., a'_t, h_2^t).$$

Therefore, we have

$$\mathbb{P}_\sigma(h_2^{t+1}|h_2^t) =$$
$$\Sigma_2 \left(\sum_{a'_1,..,a'_t,c'_1,...,c'_t} \mathbb{P}_\sigma(a'_1, .., a'_t|h_2^t)\mathbb{P}_\sigma(c'_1, ..., c'_t|h_2^t)\mathbb{P}_\sigma(\delta_{t+1}|c'_1, .., c'_t, h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|a'_1, ..., a'_t, h_2^t) \right)$$
$$= \Sigma_2\mathbb{P}_\sigma(\delta_{t+1}|h_2^t)\mathbb{P}_\sigma(\alpha_{t+1}|h_2^t).$$

It follows that the right-hand side of Equation 4 is $\mathbb{P}_\sigma(h_2^{t+1}|h_2^t)\mathbb{P}_\sigma(h_2^t) = \mathbb{P}_\sigma(h_2^{t+1})$, as required.

To complete the proof, it is enough to adapt the strategies to the information players actually have and to interpret the actions as messages sent and received (remember that Lemma 2 holds for the auxiliary game, regardless of the strategy profile). □

Lemma 9 *Consider two disjoint subsets of nodes U and V. Let $\vec{U}$ be the set of nodes $i \in \mathcal{V} \setminus U$ such that all directed paths from i to some $j \in V$, go through U.*

For all $\tau_{\vec{U}}$, $\tilde{\tau}_{\vec{U}}$, $\sigma_{-(U\cup\vec{U})}$, and for every history h_V^{T+1} of V, we have

$$\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{-(U\cup\vec{U})}}(h_V^{T+1}) = \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{-(U\cup\vec{U})}}(h_V^{T+1}).$$

Proof First, note that $\vec{U}$ is disjoint from V and that $V \subseteq V' := \mathcal{V} \setminus (U \cup \vec{U})$. Moreover, for each $i \in V'$ and each $j \in D(i)$, we have $j \notin \vec{U}$. To see this, suppose that $j \in \vec{U}$. If there exists a directed path from i to $k \in V$ that does not intersect U, then there exists a path from j to $k \in V$ that does not intersect U, a contradiction with $j \in \vec{U}$. So, either all directed paths from i to any $k \in V$ intersect U or there is no directed path from i to $k \in V$. From the definition of $\vec{U}$, it follows that $i \in \vec{U}$, a contradiction with $i \in V'$ since V' and $\vec{U}$ are disjoint.

It follows that $\cup_{i\in V'} D(i) := D(V') \subseteq V' \cup U$.

Recall that h_i^t denotes an history of messages sent and received by node i up to round t and let $h_{V'}^t := (h_i^t)_{i\in V'}$ the profile of histories of nodes in V'. We complete the proof by induction on t. For $t = 2$, we clearly have that $\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(h_{V'}^2) = \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(h_{V'}^2)$ since $D(V') \subseteq V' \cup U$ and V' is disjoint from $\vec{U}$. Assume that for $t > 2$ we have $\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(h_{V'}^t) = \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(h_{V'}^t)$. Note that the history $h_{V'}^{t+1}$ at round $t+1$ is $(h_{V'}^t, m_{V'}^{t+1}, m_{D(V')}^{t+1})$, where $m_{V'}^{t+1}$ are the messages sent by nodes in V' and $m_{D(V')}^{t+1}$ the messages received. Then,

$$\begin{aligned}
\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(h_{V'}^t, m_{V'}^{t+1}) &= \mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(h_{V'}^t)\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(m_{V'}^{t+1}|h_{V'}^t),\\
&= \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(h_{V'}^t)\sigma_{V'}(h_{V'}^t)[m_{V'}^{t+1}],\\
&= \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(h_{V'}^t)\mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(m_{V'}^{t+1}|h_{V'}^t),\\
&= \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(h_{V'}^t, m_{V'}^{t+1}).
\end{aligned}$$

Since $D(V') \subseteq V' \cup U$, we can write $m^{t+1}_{D(V')}$ as $(m^{t+1}_{D(V')\cap U}, m^{t+1}_{D(V')\cap V'})$. By definition of 0_U, the first element is the null message m_0, regardless of the history and, thus, is independent of $\tau_{\vec{U}}$. Finally, the induction hypothesis implies that $\mathbb{P}_{0_U,\tau_{\vec{U}},\sigma_{V'}}(m^{t+1}_{D(V')\cap V'}) = \mathbb{P}_{0_U,\tilde{\tau}_{\vec{U}},\sigma_{V'}}(m^{t+1}_{D(V')\cap V'})$. This completes the proof of Lemma 9. □

References

[1] Ittai Abraham, Danny Dolev, Rica Gonen and Joe Halpern, 2006. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. Proceedings of the 25th ACM Symposium on Principles of Distributed Computing, 53-62.

[2] Ittai Abraham, Danny Dolev and Joe Halpern, 2008. Lower bounds on implementing robust and resilient mediators, in TCC, Springer.

[3] Jorgen Bang-Jensen and Gregory Gutin, 2007. Digraphs Theory, Algorithms and Applications, Springer-Verlag.

[4] Imre Bàràny, 1992. Fair distribution protocols or how the players replace fortune. Mathematics of Operations Research, 17, 327-340.

[5] Amos Beimel, 2011. Secret-sharing schemes: a survey. Coding and cryptology: lecture notes in computer science, vol 6639, 11-46.

[6] Elchanan Ben-Porath, 2003. Cheap talk in games with incomplete information. Journal of Economic Theory, 108, 45-71.

[7] Elchanan Ben-Porath and Michael Kahneman, 1996. Communication in repeated games with private monitoring. Journal of Economic Theory, 70, 281–297.

[8] Benny Chor, Shafi Goldwasser, Silvio Micali and Baruch Awerbuch, 1985. Verifiable secret sharing and achieving simultaneity in the presence of faults. FOCS'85: Proceedings of the 26th Annual Symposium on Foundations of Computer Science, 383-395.

[9] Vincent P. Crawford and Joel Sobel, 1982. Strategic Information Transmission. Econometrica, 50, 1431–1451.

[10] Yvo Desmedt and Yongge Wang, 2002. Perfectly secure message transmission revisited. Lecture Notes in Computer Science, Advances in Cryptology — EUROCRYPT 2002, Vol. 2332/2002, 2002, 502-517.

[11] Danny Dolev, Cynthia Dwork, Orli Waarts and Moti Yung, 1993. Perfectly secure message transmission. Journal of the ACM, 40, 17–47.

[12] Françoise Forges, 1990. Universal mechanisms. Econometrica, 58, 1341-1364.

[13] Matthew K. Franklin and Rebecca N. Wright, 2000. Secure communication in minimal connectivity models. Journal of Cryptology, 13, 9-30.

[14] Dino Gerardi, 2004. Unmediated communication in games with complete and incomplete information. Journal of Economic Theory, 114, 104-131.

[15] S. Dov Gordon and Jonathan Katz, 2006. Rational secret sharing, revisited. In R. D. Prisco and M. Yung, editors, SCN, volume 4116 of Lecture Notes in Computer Science, pages 229-241. Springer.

[16] Joseph Y. Halpern and Vanessa Teague, 2004. Rational secret sharing and multiparty computation: extended abstract. In L. Babai, editor, STOC, pages 623-632. ACM.

[17] Kamal Jain, 2004. Security based on network topology against the wiretapping attack. IEEE Wireless Communications, 11, 69-71.

[18] Gillat Kol and Moni Naor, 2008. Cryptography and game theory: Designing protocols for exchanging information. In 5th Theory of Cryptography Conference TCC, pages 320-339.

[19] George Mailath and Larry Samuelson, 2006. Repeated Games and Reputations: Long-Run Relationships. Oxford University Press.

[20] Ueli Maurer, 1999. Information-theoretic cryptography. Advances in Cryptology — CRYPTO '99, Lecture Notes in Computer Science, Springer-Verlag, vol. 1666, 47-65.

[21] Dov Monderer and Moshe Tennenholtz, 1999. Distributed games: From mechanisms to protocols. Sixteenth National Conference on Artificial Intelligence, 32–37.

[22] Jérôme Renault and Tristan Tomala, 1998. Repeated proximity games. International Journal of Game Theory, 27, 539–559.

[23] Jérôme Renault and Tristan Tomala, 2004. Repeated proximity games. Games and Economic Behavior, 47, 124–156.

[24] Jérôme Renault and Tristan Tomala, 2008. Probabilistic reliability and privacy of communication using multicast in general neighbor networks. Journal of Cryptology, 21, 250-279.

[25] Ludovic Renou and Tristan Tomala, 2012. Mechanism design and communication networks. Theoretical Economics, 7, 489-533

[26] Adi Shamir, 1979. How to share a secret. Communications of the ACM, 22, 612-613.

[27] Jeffrey Shneidman and David C. Parkes, 2004. Specification Faithfulness in Networks with Rational Nodes. Proc. 23rd ACM Symp. on Principles of Distributed Computing (PODC'04), 88-97.

[28] Victor Shoup, 2008. A Computational Introduction to Number Theory and Algebra. Cambridge University Press.

[29] Tristan Tomala, 2011. Fault reporting in partially known networks and folk theorems. Operations Research, 59, 754-763.