A Complete Characterization of Tolerable Adversary Structures for Secure Point-to-Point Transmissions Without Feedback

Yvo Desmedt[1,3,⋆], Yongge Wang[2,⋆⋆], and Mike Burmester[3,⋆⋆⋆]

[1] University College London, UK
y.desmedt@cs.ucl.ac.uk
[2] UNC Charlotte, USA
yonwang@uncc.edu
[3] Florida State University, USA
burmester@cs.fsu.edu

Abstract. Problems of unconditionally secure communication have been studied extensively in network models. Dolev-Dwork-Waarts-Yung considered the Byzantine threats model, in which the adversary can only take over a number of nodes bounded by a threshold. They studied two cases:

1. all communication links (edges in the graph) are two-way communication,
2. all communication links are one-way communication, and there is no feedback.

The node sets that the adversary can take over was generalized by Hirt-Maurer to an adversary structure. At PODC 2002, Kumar-Goundan-Srinathan-Rangan generalized Dolev-Dwork-Waarts-Yung's first scenario to the case of a general adversary structure. In this paper we generalize Dolev-Dwork-Waarts-Yung's second scenario to the case of a general adversary structure. As in Dolev-Dwork-Waarts-Yung, our work relies on the use of secret sharing.

Keywords: Network security, Byzantine threats, Secret Sharing, adversary structure, unconditional security.

1 Introduction

Originally work on secure distributed computation (see, e.g., [1,3]) assumed that the parties were connected by a complete network of reliable and private point-to-point communication channels, with an adversary that could control up to k (Byzantine) nodes.

In practical networked environments it is often the case that two parties are not connected with a private and authenticated channel. They then need to use intermediate nodes to help them to carry out secure communication and secure multiparty computations. In this case, it is important to design secure communication protocols to achieve

⋆ A part of this work has been funded by CCR-0209092. The author is BT Professor of Information Security.

⋆⋆ This work was supported, in part, by funds provided by The University of North Carolina at Charlotte.

⋆⋆⋆ A part of this work has been funded by CCR-0209092.

X. Deng and D. Du (Eds.): ISAAC 2005, LNCS 3827, pp. 277–287, 2005.

participant cooperation in the presence of faults. Dolev, Dwork, Waarts, and Yung [6] initiated the study of reducing the requirements for network connectivity in secure multiparty computation by providing protocols that achieve private and reliable communication. In the case of k Byzantine faults, they studied the cases:

1. all communication links (edges in the graph) are two-way communication. Reliable and private communication is achievable if and only if the communication network is $2k + 1$ connected.
2. all communication links are one-way communication, and there is no feedback. Then $3k + 1$ connectivity is necessary and sufficient for reliable and private communications.

In 2002, Desmedt and Wang [4] extended this to the case when there are feedback channels.

The Byzantine faults model typically addresses threat scenarios in which the faults are *independent*. This assumption is unrealistic when dealing with computer viruses, such as the ILOVEYOU [11] virus and the Internet virus/worm [7] that only spread to Windows, respectively Unix. A hacker who can exploit a weakness in one platform, can with almost the same ease attack many computers, if not all, on that same platform. There is therefore a need to design a model in which nodes on the same platform can be distinguished. One such model, the *adversary structure* model, was proposed by Hirt and Maurer [10] for secure multiparty computation (for an earlier version see [9]). In this model the adversary is characterized by a structure (or family) of subsets of the set of parties, which are the sets that the adversary can corrupt. Hirt and Maurer gave necessary and sufficient conditions for secure multiparty computation in this adversary model. Similar to the classical results for multiparty computation, Hirt and Maurer assumed that communication networks are complete.

As in Dolev, Dwork, Waarts, and Yung [6], in this paper, we study the problem of reducing the requirements for network connectivity in secure multiparty computation in the sense of Hirt and Maurer [10] under the adversary structure model. We give necessary and sufficient conditions on the communication network, with respect to a given adversary structure $\mathcal{Z}$, such that we have reliable and private point-to-point communication. It should be noted that results for not necessarily complete bi-direction networks have been obtained by Kumar-Goundan-Srinathan-Rangan in [13]. They used an induction argument to prove the sufficient condition. There are two essential differences between our results and the results in [13].

– The secure message transmission protocols in [13] only apply to networks in which all links are two-way. That is, one can send messages in two directions. Our protocols work for networks that are not necessarily complete and the links are one-way. That is, one may send message from one direction though not the other direction.
– Our protocols for secure transmissions are slightly more efficient than these in [13] due to the induction processes in [13]. Specifically, our protocols will use *minimal* $\mathcal{Z}$-connected path-sets. We shall find a bound on the number of paths in such path-sets, and show that it is sharp. Although we focus on point-to-point networks, our results can easily be extended to broadcast networks.

The organization of this paper has as follows. In Section 2, we describe our model and give definitions. In Section 3 we find necessary and sufficient conditions for secure communication in our adversary model and propose secure transmission protocols. In Section 4 we propose bounds.

2 Model and Background

2.1 Threshold Secret Sharing Schemes

Let $\mathbf{F}$ be a finite field, and let k, n be integers such that $0 \leq k < n$. A $(k+1)$-out-of-n *secret sharing scheme* is a probabilistic function S: $\mathbf{F} \rightarrow \mathbf{F}^n$ with the property that: for any $M \in \mathbf{F}$ and $\mathrm{S}(M) = (v_1, \ldots, v_n)$, no information about M can be inferred from any k entries of $(v_1, \ldots, v_n)$, whereas M can be recovered from any $k+1$ entries of $(v_1, \ldots, v_n)$.

2.2 The Adversary

We employ an unconditional setting. That is, the adversary has unlimited resources. The adversary is characterized by an adversary structure $\mathcal{Z}$ [10] that consists of all sets of parties that the adversary can corrupt. Formally, let $\mathcal{P}$ be a party set. A subset $\Gamma_{\mathcal{P}}$ of the power set $2^{\mathcal{P}}$ of $\mathcal{P}$ is called an access structure on $\mathcal{P}$ [12]. It is monotone if and only if $\emptyset \notin \Gamma_{\mathcal{P}}$ and supersets of elements $\Gamma_{\mathcal{P}}$ also belong to $\Gamma_{\mathcal{P}}$, i.e., we require that if $A \in \Gamma_{\mathcal{P}}$ and $A \subseteq A' \subseteq \mathcal{P}$, then $A' \in \Gamma_{\mathcal{P}}$. Let $\mathcal{Z}_{\mathcal{P}} = \Gamma_{\mathcal{P}}$. We call $\mathcal{Z}_{\mathcal{P}} \subset 2^{\mathcal{P}}$, or simply $\mathcal{Z}$ when there is no confusion, an adversary structure on $\mathcal{P}$ if its complement, i.e., $\mathcal{Z}_{\mathcal{P}}^c = 2^{\mathcal{P}} \setminus \mathcal{Z}_{\mathcal{P}}$ is a monotone access structure.

If $\mathcal{Z}_1$ and $\mathcal{Z}_2$ are adversary structures for $\mathcal{P}$, then $\mathcal{Z}_1 + \mathcal{Z}_2 = \{Z_1 \cup Z_2 : Z_1 \in \mathcal{Z}_1, Z_2 \in \mathcal{Z}_2\}$ is also an adversary structure for $\mathcal{P}$. $2\mathcal{Z}$ and $3\mathcal{Z}$ are the adversary structures $\mathcal{Z} + \mathcal{Z}$ and $\mathcal{Z} + \mathcal{Z} + \mathcal{Z}$ respectively. Finally, a set of parties $Z \in \mathcal{Z}$ is *maximal* if: $Z' \supset Z \Rightarrow Z' \notin \mathcal{Z}$. In the remaining part of the paper, we will use $2\mathcal{Z}$ and $3\mathcal{Z}$ to denote the adversary structures $\mathcal{Z} + \mathcal{Z}$ and $\mathcal{Z} + \mathcal{Z} + \mathcal{Z}$ respectively.

We consider two kinds of adversary: a *passive* and an *active* adversary. A passive adversary (or gossiper adversary) can only read the traffic (the variables in the view) of the parties in Z. An active adversary has unlimited computational power and can read the traffic of Z and also control the parties in Z. Both kinds of adversary are assumed to know the complete protocol specification, message space, and the complete structure of the graph. In this paper, we shall not consider dynamic adversaries who can change the parties they corrupt from round to round, but only static adversaries. That is, the adversary selects which set of parties $Z \in \mathcal{Z}$ to corrupt, before the start of the protocol.

2.3 The Communication Network

We model the communication network by a directed graph $G = G(V, E)$ whose nodes are the parties and whose edges are point-to-point reliable and private communication channels.

2.4 Message Transmission Protocols

Let π be a message transmission protocol, let A be the sender and B the receiver, and let $\mathcal{Z}$ be an adversary structure. The sender A selects a message M^A drawn from a message space $\mathcal{M}$ with a certain probability distribution. At the beginning of the protocol, the adversary flips coins and chooses a set $Z \in \mathcal{Z}$ of nodes to corrupt. At the end of protocol π, the receiver B outputs a message $M^B \in \mathcal{M}$. We will assume that the message space $\mathcal{M}$ is a subset of a finite field $\mathbf{F}$ (our results easily extend to message spaces of tuples over $\mathbf{F}$).

For any message transmission protocol we denote by adv the adversary's view of the execution of the protocol and by $adv(M, r)$ the view when $M^A = M$ and when the sequence of coin flips used by the adversary is r.

Definition 1. *Let π be transmission protocol, let M^A the message selected by A and M^B the message output by B, let $\mathcal{Z}$ be an adversary structure.*

1. *We say that π is* $\mathcal{Z}$-reliable *if B outputs $M^B = M^A$ with probability* 1. *(The probability is taken over the choices of M^A and the coin flips of all nodes.)*
2. *We say that π is perfectly* $\mathcal{Z}$-private *if for any two messages M_0, M_1, and for any coin tosses r, we have $\Pr[adv(M_0, r) = c] = \Pr[adv(M_1, r) = c]$. The probabilities are taken over the coin flips of the honest parties.*
3. *We say that π is perfectly* $\mathcal{Z}$-secure *if it is* $\mathcal{Z}$-reliable *and perfectly* $\mathcal{Z}$-private.

2.5 Connectivity

Definition 2. *Let $G(V, E)$ be a directed graph, A, B be nodes in $G(V, E)$, and $\mathcal{Z}$ be a an adversary structure on $V \setminus \{A, B\}$.*

- *A, B are* $\mathcal{Z}$-separable *in G, if there is a set $Z \in \mathcal{Z}$ such that all paths from A to B go through at least one node in Z. We say that Z separates A and B.*
- *A, B are* $(\mathcal{Z}+1)$-connected *if they are not $\mathcal{Z}$-separable in G.*

Observe that if $(A, B) \in E$, then A, B are $(\mathcal{Z}+1)$-*connected* for any $\mathcal{Z}$ on $V \setminus \{A, B\}$. The following result will be needed in our later discussions.

Theorem 1. *Let $G = G(V, E)$ be a directed graph, A, B be nodes in G, and $\mathcal{Z}_1, \mathcal{Z}_2$ be adversary structures on $V \setminus \{A, B\}$. Then A, B are $(\mathcal{Z}_1 + \mathcal{Z}_2 + 1)$-connected if, and only if: for all sets $Z_1 \in \mathcal{Z}_1$ there is a set S_{Z_1} of paths between A and B such that,*

- *the paths in S_{Z_1} are free from nodes of Z_1,*
- *for every $Z_2 \in \mathcal{Z}_2$ there is at least one path in S_{Z_1} that is free from nodes of Z_2.*

Proof. First consider the case when A, B are $(\mathcal{Z}_1 + \mathcal{Z}_2 + 1)$-connected. We prove that the conditions are satisfied. For any set $Z_1 \in \mathcal{Z}_1$, let S_{Z_1} be the set of all paths from A to B free from nodes of Z_1. Assume that there is a set $Z_2 \in \mathcal{Z}_2$ such that all paths of S_{Z_1} go through Z_2. Then $Z_1 \cup Z_2$ separates A, B in G. That is, A, B are $(\mathcal{Z}_1 + \mathcal{Z}_2)$-separable in G. This is a contradiction.

For the converse observe that the conditions on the paths S_{Z_1} make it impossible to have a set $Z_2 \in \mathcal{Z}_2$ such that $Z = Z_1 \cup Z_2$ separates A, B. Indeed if there where such a set Z separating A, B then there would be no path in S_{Z_1} free of Z_1 and Z_2.

3 Secure Message Transmissions

We start by discussing message transmissions that tolerate passive adversaries. First we briefly describe the intuition behind our protocols by observing that, whereas in the Byzantine threats model the sender and receiver use vertex disjoint paths, for general adversary structures this is not necessarily the case.

Theorem 2. *Let $G = G(V, E)$ be a directed graph, A, B be two nodes in G, and $\mathcal{Z}$ be an adversary structure on $V \setminus \{A, B\}$. Suppose that the adversary is passive.*

1. *We have polynomial time (in the graph size) $\mathcal{Z}$-reliable message transmission from A to B if, and only if, A, B are $(\{\emptyset\} + 1)$-connected in G.*
2. *We have polynomial time (in the graph size) perfectly $\mathcal{Z}$-secure message transmission from A to B if, and only if, A, B are $(\mathcal{Z} + 1)$-connected in G.*

Proof. Result 1 is straightforward. For Result 2, first observe that if A, B are $\mathcal{Z}$-separable in G, then there is a set $Z \in \mathcal{Z}$ such that all paths from A go through Z. Therefore $\mathcal{Z}$-private message transmission from A to B is impossible. Next suppose that A, B are $(\mathcal{Z} + 1)$-connected in G. In the following we describe a polynomial time (in the graph size) protocol for A to securely send a message M^A to B (our protocol is similar to a protocol in [8, Lemma 2]).

1. Let $G'(V'E')$ be the maximum subgraph of $G(V, E)$ such that for each node $v \in V'$, there is a direct path from v to B.
2. For each edge $(u, v) \in E'$, u chooses a random message (group elements chosen according to the uniform distribution) $r_{u,v}$ and sends it to the node v.
3. Every node computes the sum of messages it has received and substracts the sum of messages it has sent out. If the node is the actual sender A, then it adds to this total the messsage M^A. Call this sum the "final result" for this node.
4. Each final result from Step 3, except for the final result held by the actual receiver B, is propagated by the nodes openly to the receiver B through a series of transmissions. The sum of all final results, including the final result held by the receiver B, is the message M^B.

Using a similar proof to that in [8, Lemma 2], we can show that the above protocol is a $\mathcal{Z}$-secure message transmission from A to B if A, B are $(\mathcal{Z} + 1)$-connected in G.

Next we consider secure message transmission protocols for active adversaries.

Theorem 3. *Let $G = G(V, E)$ be a directed graph, A, B be nodes in G, and $\mathcal{Z}$ be an adversary structure on $V \setminus \{A, B\}$. We have $\mathcal{Z}$-reliable message transmission from A to B if, and only if, A, B are $(2\mathcal{Z} + 1)$-connected in G.*

Proof. First assume that A, B are $(2\mathcal{Z} + 1)$-connected in G, and let S be the set of all directed paths from A to B. The paths in S will be used by the sender A to transmit messages to B. Let M^A be the message that A wants to send to B. For each path $p \in S$, A sends M^A to B over p. At the end of the protocol, B receives M_p^B through each path $p \in S$. Then by using Theorem 1, B finds a node set $Z_1 \in \mathcal{Z}$ whose path set S_{Z_1} is such that the same message is received on all its paths. Let M^B be this message. It is sufficient to show that $M^B = M^A$. Suppose that the adversary selects $Z_2 \in \mathcal{Z}$. We have:

- If $Z_1 \cap Z_2 = \emptyset$ then $M^A = M^B$.
- If $Z_1 \cap Z_2 \neq \emptyset$ then by Theorem 1, since A, B are $(2\mathcal{Z}+1)$-connected, there will be a path $p_0 \in S_{Z_1}$ free from nodes of Z_2. On this path $M^B_{p_0} = M^A$. Since B receives the same message from all paths in S_{Z_1}, we must have $M^A = M^B_{p_0} = M^B$.

It follows that B can reliably recover the message M^A.

Next assume that A, B are not $(2\mathcal{Z}+1)$-connected in G. That is, there are sets $Z_1, Z_2 \in \mathcal{Z}$ such that all directed paths from A to B passes through some nodes in $Z_1 \cup Z_2$. Let M_0 be the message that A transmits. The adversary will attempt to maintain a simulation of the possible behavior of A by executing the message transmission protocol for some other message M_1. The strategy of the adversary is to flip a coin and then, depending on the outcome, decide which set of Z_1 or Z_2 to control (our model is not dynamic, so the selection is done before the protocol starts). Let Z_b be the chosen set. In each execution step of the transmission protocol, the adversary causes each node in Z_b to follow the protocol as if the transmitted message were M_1. Since B does not know whether $b = 1$ or $b = 2$, with probability better than 1/2, at the end of the protocol B cannot decide whether A has transmitted M_0 or M_1 with probability better than 1/2.

Theorem 4. *Let $G = G(V, E)$ be a directed graph, A, B be nodes in G, and $\mathcal{Z}$ be an adversary structure on $V \setminus \{A, B\}$. If there are no directed paths from B to A, then we have perfectly $\mathcal{Z}$-secure message transmission from A to B if and only if, A, B are $(3\mathcal{Z}+1)$-connected in G.*

Proof. First we show that the condition is sufficient. We shall describe a message transmission protocol π that is perfectly $\mathcal{Z}$-secure. Let $\mathcal{Z} = \{Z_1, \ldots, Z_t\}$ be the adversary structure (to make our protocol more efficient, we could take a list of maximal adversary sets) and let M^A be the message that A wants to send to B. The sender A first uses a t-out-of-t secret sharing scheme to get shares $(s^A_1, \ldots, s^A_t)$ of the message M^A. For each $i \leq t$, A reliably sends s^A_i to B via the reduced graph $G_{V \setminus Z_i} = G_{V \setminus Z_i}(V \setminus Z_i, E_{V \setminus Z_i})$, where $E_{V \setminus Z_i} = E \setminus \{(u, v) : u \in Z_i \text{ or } v \in Z_i\}$. For each $i \leq t$, B reliably receives s^B_i on the reduced graph $G_{V \setminus Z_i}$, and hence recovers the message M^B from the shares $(s^B_1, \ldots, s^B_t)$. Now assume that the adversary controls all nodes in Z_{i_0}. Then the adversary will learn no information about M^A from $s^A_{i_0}$. Therefore the transmission protocol π is perfectly $\mathcal{Z}$-private. It remains to show that π is $\mathcal{Z}$-reliable. Since A, B are $(3\mathcal{Z}+1)$-connected, it is straightforward to see that for each $Z_i \in \mathcal{Z}$, the reduced graph $G_{V \setminus Z_i}$ is $(2\mathcal{Z}+1)$-connected. From Theorem 3 we get that B receives reliably all the shares $(s^A_1, \ldots, s^A_t)$. It follows that the protocol is perfectly $\mathcal{Z}$-secure.

Next we show that the condition is necessary. Assume that A, B are not $(3\mathcal{Z}+1)$-connected. Then there are sets $Z_1, Z_2, Z_3 \in \mathcal{Z}$ such that all directed paths from A to B pass through some nodes in $Z_1 \cup Z_2 \cup Z_3$. Let M_0 be the message that A transmits. The adversary will attempt to maintain a simulation of the possible behavior of A by executing the transmission protocol π for some other message M_1. The strategy of the adversary is to flip coins and then, depending on the outcome, decide which set of Z_1, Z_2 or Z_3 to control. Let Z_b be the chosen set. In each execution step of the transmission protocol, the adversary causes each node in Z_b to follow the protocol π as if the protocol were transmitting the message M_1. At the end of the protocol, the view of B could be divided into three parts $view_{Z_1}, view_{Z_2}$, and $view_{Z_3}$, where $view_{Z_i}$ $(i = 1, 2, 3)$

consists of all the information that the nodes in Z_i have learned. Since neither A nor B knows whether $b = 1$, $b = 2$ or $b = 3$, and since π is a perfectly private message transmission protocol, M^A cannot be recovered from any single $view_{Z_i}$. Since π is a reliable message transmission protocol and the adversary controls one of Z_1, Z_2, or Z_3, B should be able to recover the secret message from two of these three views. That is, if we regard $(view_{Z_1}, view_{Z_2}, view_{Z_3})$ as shares of M^A in a 2-out-of-3 secret sharing scheme, then this scheme should be able to detect and simultaneously correct one error. However 2-out-of-3 secret sharing schemes can only detect one error, and cannot correct any errors (see, e.g., [14]). So we get a contradiction. This proves that there is no perfectly secure message transmission protocol from A to B when A, B are only $3\mathcal{Z}$-connected.

4 Bounds and Other Properties

4.1 Introduction

A *threshold adversary structure* $\mathcal{Z}$ is a structure whose maximal sets Z have cardinality bounded by a threshold k. An example of such a structure is the classical Byzantine faults structure. Evidently if one restricts oneself to threshold adversary structures, then there is a description of the adversary structure which is logarithmic in the number of elements in it. In this context, it should be noted that the "polynomial" algorithms for multiparty computation in Hirt and Maurer [10] are polynomial in the size of the adversary access structure which is generally exponential in the size of the network. Hirt and Maurer do not study the problem for restricted access structures.

For a threshold adversary structure $\mathcal{Z}$ we get $\mathcal{Z}$-tolerable communication in both the passive and active case [5,6] via path-sets S for which,

- the paths are vertex-disjoint, and so
- the number of paths is polynomially bounded (in the size of the graph).

The goal of this section is to show that both properties are false in the general adversary case, when limiting the adversary structures. Our results do *not* rely on unproven assumptions.

We shall use a family of specific adversary structures $\mathcal{Z}$ that was informally introduced in [2]. This structure consists of sets Z of nodes of a colored graph that have at most k colors (e.g. computers running the same/different operating system are respectively colored with the same/different color). In other words, the adversary can corrupt any set of nodes provided it has no more than k colors. Evidently the number of nodes in Z can be much larger than k (if many nodes have the same color). We now define this structure formally.

Definition 3. *A colored graph is a tuple $G = G(V, E, C, f)$, with V the node set, E the edge set, C the color set, and f a map from V onto C. The structure*

$$\mathcal{Z}_{C,k} = \{Z \mid Z \subset V \text{ and } |f(Z)| \leq k\}.$$

is called a k-color adversary structure.

Note that if all nodes of the colored graph have different colors then we have the classical Byzantine faults structure, with no colors.

4.2 Bounds

To study the properties of path-sets that we shall use for $\mathcal{Z}$-tolerable communication we define the following. (This definition can also be used for general adversary structures, not based on colors.)

Definition 4. *Let $G(V, E)$ be a directed graph, A, B be nodes in G, S be a set of simple paths in G between A and B, and G_S be the graph obtained by removing all nodes and edges of G not in S. Let $\mathcal{Z}$ be an adversary structure. We say that S is a minimal $(\mathcal{Z}+1)$-connected path-set from A to B in G, if*

1. *A and B are $(\mathcal{Z}+1)$-connected in G_S, and*
2. *for each path $p \in S$, A and B are $\mathcal{Z}$-separable in $G_{S\setminus\{p\}}$.*

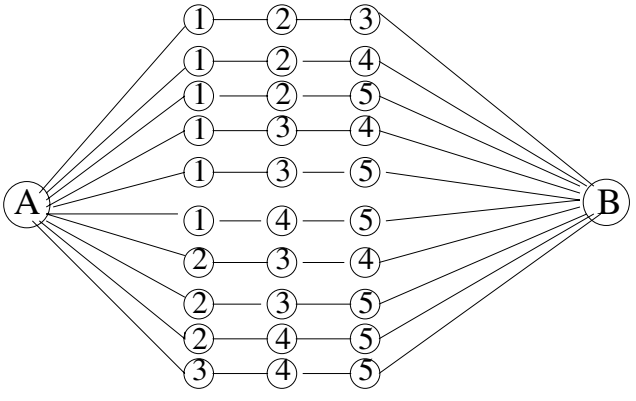

Fig. 1. A minimal $(\mathcal{Z}_{C,k}+1)$-connected path-set of a graph with 5 colors for a 2-color adversary structure ($c = 5$, $k = 2$)

Lemma 1. *Let $G = G(V, E, C, f)$ be a colored graph, S be a minimal $(\mathcal{Z}_{C,k}+1)$-connected path-set between A and B, and c be the number of colors in the graph. Then $|S| \leq \binom{c}{k}$. This bound is tight.*

Proof. We apply Theorem 1 with $\mathcal{Z}_1 = \{\emptyset\}$ and $\mathcal{Z}_2 = \mathcal{Z}_{C,k}$, and construct a minimal path-set, starting with $S = \emptyset$. For every maximal $Z_2 \in \mathcal{Z}_2$, find a path free from nodes of Z_2, label it Z_2, and add it to S. If there is already such a path in S, then just add the extra label Z_2 to this path.

Note that non-maximal sets do not add extra paths. Indeed, if after running this construction one takes a set $Z_2' \subset Z_2$, then the path with label Z_2 will also receive the Z_2' label. Since all the different sets $Z_2 \in \mathcal{Z}_2$ have been considered, the set S will be $(\mathcal{Z}_{C,k}+1)$-connected (by Theorem 1).

Now note that a maximal set Z_2 must have k colors and Z_2 must contain all nodes with these colors. So there are $\binom{c}{k}$ maximal sets in $\mathcal{Z}_2$. Combining this with the previous argument we get $|S| \leq \binom{c}{k}$.

To show that the bound is tight we describe an appropriate colored graph $G = G(V, E, C, f)$ and a minimal path-set S. Figure 1 illustrates our proof for $c = 5$ and $k = 2$.

Let $c \geq k+1$. We construct V and E as follows. Start with $V = \{A, B\}$ and $E = \emptyset$. Then add $\binom{c}{k}$ node-disjoint paths connecting A to B as follows:

For i *from* 1 *till* $\binom{c}{k}$ *do*:

Step 1 Select a not yet selected subset $C' \subset C$ of k colors.
Step 2 Add $c - k$ new nodes $v_{i,1}, \ldots, v_{i,c-k}$ to V.
Step 3 Color these $c - k$ new nodes with the $C \setminus C'$ different colors, in such a way that each gets a different color.
Step 4 Add the edges $(A, v_{i,1})$, $(v_{i,1}, v_{i,2})$, …, $(v_{i,c-k-1}, v_{i,c-k})$, $(v_{i,c-k}, B)$ to E.

end-for-loop.

It is easy to verify that we get a minimal $(\mathcal{Z}_{C,k} + 1)$-connected path-set from A to B. Indeed, if one picks a maximal element from $\mathcal{Z}_{C,k}$ and removes all its nodes (which is equivalent to picking any k colors and removing all nodes with these colors) one is left with just one path. So, A and B are $(\mathcal{Z}_{C,k} + 1)$-connected. This also implies that if one removes a single path from this path-set then A, B become $\mathcal{Z}_{C,k}$-separable.

Remark 1. Though the minimal path-set in Lemma 1 is super-polynomial in k, it is still polynomial in the graph size. The construction in the proof of Lemma 1 will be used to prove Theorem 5.

Theorem 5. *There are directed graphs $G(V, E)$ and adversary structures $\mathcal{Z}$ for which the number of paths in a minimal $(\mathcal{Z} + 1)$-connected path-set from A to B is super-polynomial in the size of the graph.*

Proof. We modify the colored graph $G(V, E, C, f)$ constructed in Lemma 1 to get a graph $G'(V', E')$ as follows. Take $|C| = 2k + 1$ (so $c - k = k + 1$) and construct G' with $|E'| = |E|$ and $|V'| = (k + 1)^2 + 2$. Start from G. For each path in this graph reorder linearly the nodes accordingly to their colors (if the colors are labeled $1, 2, \ldots$ put the nodes with the smallest color label the closest to A). Make the graph directed, i.e. the paths go from A to B. Write the nodes and their colors in a $\binom{c}{k} \times (k+1)$ matrix $T = [(v_{i,j}, f(v_{i,j})]$.

Now we explain how V' is constructed. In each column of T collapse the nodes with the same color: that is, if $f(v_{i,j}) = f(v_{i',j})$ then $v'_{i,j} = v'_{i',j}$. From elementary combinatorics we get that the number of different nodes that remain in each column is $k + 1$. So, we have a total of $(k + 1)^2 + 2$ different nodes including A and B. It is now easy to verify that the size of this path-set is super-polynomial.

Observe that the paths of the minimal path-set in Lemma 1 are *not* node-disjoint.

Remark 2. Though the number of paths in a minimal $(\mathcal{Z}+1)$-connected path-set for the directed graph $G(V, E)$ in Theorem 5 is super-polynomial in the graph size, one may still be able to design polynomial time algorithms (in the size of the graph) for perfect

secure communication. Indeed, for the case of privacy against a passive adversary, we have an algorithm that runs in polynomial time in the graph size, which is more efficient than in polynomial time in the graph and adversary size (see Theorem 2). It is an **open** problem whether there exists a graph such that all perfect secure communication protocols are super-polynomial in the graph size.

More details for colored graphs are given in the full version of this paper.

4.3 Other Results for Color Adversary Structures

Since color adversary structures are of interest on their own (to model platform dependent attacks) [2], we prove the following result:

Theorem 6. *Let $G = G(V, E, C, f)$ be a colored graph which is $(\mathcal{Z}_{C,k}+1)$-connected. If the number of colors is minimal then the paths in a minimal path-set are node-disjoint and each path is monochrome (all nodes on one path have the same color).*

Proof. Using Lemma 1 we see that the number of colors must be at least $k + 1$. We now prove that we can have $k + 1$ colors and that there is a minimal path-set with $k + 1$ monochrome paths. Apply Theorem 1 with $\mathcal{Z}_1 = \{\emptyset\}$ and $\mathcal{Z}_2 = \mathcal{Z}_{C,k}$ to construct a minimal path-set. Start with $S = \emptyset$. Let Z_2 be a maximal subset of $\mathcal{Z}_2$. Obviously Z_2 contains nodes of k different colors. Moreover, due to its maximality, it contains all nodes that have these colors. Then by Theorem 1 there must be at least one path which is free from the nodes of Z_2, that is, free from the k colors of Z_2. Since there are only $k + 1$ colors, this path must be monochrome. Add one of these paths to S. Continue with a new maximal $Z_2' \in \mathcal{Z}_2$ with at least one color different from Z_2 until all k-color sets are exhausted. In this way we get a minimal path-set consisting of $k + 1$ paths, one for each color. Clearly these paths are disjoint.

References

1. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computing. In: *Proc. ACM STOC '88*, pages 1–10.
2. M. Burmester and Y. G. Desmedt. Is hierarchical public-key certification the next target for hackers? *Communications of the ACM*, 47(8), pp. 68–74, August 2004.
3. D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols. In *Proc. ACM STOC*, pp. 11–19, May 2–4, 1988.
4. Y. Desmedt and Y. Wang. Perfectly secure message transmission revisited. In: *Proc. Eurocrypt '02*, Lecture Notes in Computer Science 2332, Springer-Verlag, pp. 502–517, 2002.
5. D. Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3, pp. 14–30, 1982.
6. D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *Journal of the ACM*, 40(1), pp. 17–47, January 1993.
7. M. W. Eichin and J. A. Rochlis. With microscope and tweezers: an analysis of the Internet virus of November 1988. In *Proc. IEEE Sym. on Security and Privacy*, pp. 326–343, 1989.
8. M. Franklin and M. Yung. Secure hypergraphs: privacy from partial broadcast. In: *Proc. ACM STOC '95*, pages 36–44, ACM Press, 1995.
9. M. Hirt and U. Maurer. Complete Characterization of Adversaries Tolerable in Secure Multi-Party Computation. In *Proc. of the 16th ACM PODC*, pp. 25–34, August, 1997.

10. M. Hirt and U. Maurer. Player Simulation and General Adversary Structures in Perfect Multiparty Computation. Journal of Cryptology 13(1): 31-60 (2000)
11. 'ILOVEYOU' computer bug bites hard, spreads fast. `http://www.cnn.com/2000/TECH/computing/05/04/iloveyou.01/index.html`, May 4, 2000.
12. M. Ito, A. Saito and T. Nishizeki. Secret sharing schemes realizing general access structures. *Proc. IEEE Global Telecommunications Conf., Globecom'87*, pp. 99–102, 1987.
13. M. Kumar, P. Goundan, K. Srinathan, and C. Rangan. On perfectly secure communication over arbitrary networks. Proc. of ACM PODC 2002, pages 193–202.
14. F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes*. North-Holland Publishing Company, 1978.