J. Cryptology (2001) 14: 121–135
DOI: 10.1007/s00145-001-0002-y

Journal of
CRYPTOLOGY

Secure Communication in Multicast Channels: The Answer to Franklin and Wright's Question*

Yongge Wang
Certicom Research, Certicom Corp., 5520 Explorer Dr., 4th floor,
Mississauga, Ontario, Canada L4W 5L1
ywang@certicom.com

Yvo Desmedt
Department of Computer Science,
Florida State University, Tallahassee, FL 32306-4530, U.S.A.
desmedt@cs.fsu.edu
and
Information Security Group,
Royal Holloway—University of London,
Egham, Surrey TW20 0EX, England

Communicated by Cynthia Dwork

Received June 1999 and revised September 2000
Online publication 9 March 2001

Abstract. Problems of secure communication and computation have been studied extensively in network models. Goldreich et al., Franklin and Yung, and Franklin and Wright have initiated the study of secure communication and secure computation in multirecipient (multicast) models. A "multicast channel" (such as ethernet) enables one processor to send the same message—simultaneously and privately—to a fixed subset of processors. In their recent paper, Franklin and Wright have shown that if there are n multicast lines between a sender and a receiver and there are at most t malicious (Byzantine style) processors, then the condition $n > t$ is necessary and sufficient for achieving efficient probabilistically reliable and probabilistically private communication. They also showed that if $n > \lceil 3t/2 \rceil$, then there is an efficient protocol to achieve probabilistically reliable and perfectly private communication. They left open the question whether there exists an efficient protocol to achieve probabilistically reliable and perfectly private communication when $\lceil 3t/2 \rceil \geq n > t$. In this paper, by using a different authentication scheme, we answer this question affirmatively and study related problems.

Key words. Network security, Privacy, Perfect secrecy, Reliability.

[*] Research supported by DARPA F30602-97-1-0205. A preliminary version of this paper has appeared in *Advances in Cryptology – Eurocrypt '99 Proceedings*, Lecture Notes in Computer Science 1592, Springer-Verlag, Berlin, pp. 443–455.

1. Introduction

If two parties are connected by a private and authenticated channel, then secure communication between them is guaranteed. However, in most cases, many parties are only indirectly connected, as elements of an incomplete network of private and authenticated channels. In other words they need to use intermediate or internal nodes. Achieving participants' cooperation in the presence of faults is a major problem in distributed networks. The interplay of network connectivity and secure communication has been studied extensively (see, e.g., [1], [3], [5], [6], and [12]). For example, Dolev [5] and Dolev et al. [6] showed that, in the case of t Byzantine faults, reliable communication is achievable only if the system's network is $2t + 1$ connected. Hadzilacos [12] has shown that connectivity $t + 1$ is required to achieve reliable communication in the presence of t faulty participants even if those faults are not malicious.

Goldreich et al. [11], Franklin and Yung [8], and Franklin and Wright [7] have initiated the study of secure communication and secure computation in *multirecipient* (*multicast*) models. A "multicast channel" (such as ethernet) enables one participant to send the same message—simultaneously and privately—to a fixed subset of participants. Franklin and Yung [8] have given a necessary and sufficient condition for individuals to exchange private messages in multicast models in the presence of passive adversaries (passive gossipers). For the case of active Byzantine adversaries, many results have been presented by Franklin and Wright [7]. Note that Goldreich et al. [11] have also studied fault-tolerant computation in the public multicast model (which can be thought of as the largest possible multirecipient channels) in the presence of active Byzantine adversaries. Specifically, Goldreich et al. [11] have made an investigation of general fault-tolerant distributed computation in the full-information model. In the full information model no restrictions are made on the computational power of the faulty parties or the information available to them. (Namely, the faulty players may be infinitely powerful and there are no private channels connecting pairs of honest players). In particular, they present efficient two-party protocols for fault-tolerant computation of any bivariate function.

There are many examples of multicast channels. A simple example is a local area network like an ethernet bus or a token ring. Another example is a shared cryptographic key. By publishing an encrypted message, a participant initiates a multicast to the subset of participants that are able to decrypt it.

We abstract away the concrete network structures and consider multicast graphs. Specifically, a multicast graph is a graph $G(V, E)$. A vertex $A \in V$ is called a neighbor of another vertex $B \in V$ if there there is an edge $(A, B) \in E$. In a multicast graph, we assume that any message sent by a node A will be received identically by all its neighbors, whether or not A is faulty, and all parties outside of A's neighbor learn nothing about the content of the message. These neighbor networks have been studied by Franklin and Yung in [8]. They have also studied the more general notion of hypergraphs, which we do not need.

As Franklin and Wright [7] have pointed out, unlike in the simple channel model, it is not possible to apply protocols directly over multicast lines to disjoint paths in a general multicast graph, since disjoint paths may have common neighbors. Franklin and Wright have shown that in certain cases the change from a simple channel to a multicast channel hurts the adversary more than it helps, because the adversary suffers from the restriction that an incorrect transmission from a faulty processor will always be received identically by all of its neighbors.

Franklin and Wright [7] showed that if there are n multicast lines (that is, n paths with disjoint neighborhoods) between a sender and a receiver and there are at most t malicious (Byzantine style) processors, then the condition $n > t$ is necessary and sufficient for achieving efficient probabilistically reliable and probabilistically private communication. They also showed that there is an efficient protocol to achieve probabilistically reliable and perfectly private communication when $n > \lceil 3t/2 \rceil$, and there is an exponential bit complexity protocol for achieving probabilistically reliable and perfectly private communication when $\lceil 3t/2 \rceil \geq n > t$. However, they left open the question whether there exists an efficient protocol to achieve probabilistically reliable and perfectly private communication when $\lceil 3t/2 \rceil \geq n > t$. In this paper, using a different authentication scheme, we answer this question affirmatively and study related problems: the use of multicast line protocols when the lines are embedded in a more general multicast graph. There are limits to what we can expect to achieve in the most general case. We also show that it is **NP**-complete to decide whether a multicast graph has n disjoint multicast lines (that is, n paths with disjoint neighborhoods).

Note that, similar to Franklin and Wright [7], we only consider the scenario when the underlying graph is known to all nodes. For the scenario that the graph is unknown, the protocols may be completely different, see [2].

We present our model in Section 2. We review the relevant result of [7] in Section 3. In Section 4 we present a solution to the Franklin–Wright open question. In Section 5 we consider the applicability of multicast line protocols to general multicast graphs.

2. Model

Throughout this paper, n denotes the number of multicast lines and t denotes the number of faults under the control of the adversary. We write $|S|$ to denote the number of elements in the set S. We write $x \in_R S$ to indicate that x is chosen with respect to the uniform distribution on S. Let $\mathbf{F}$ be a finite field, and let $a, b, M \in \mathbf{F}$. We define auth $(M, a, b) := aM + b$ (following [7], [10], [13], and [14]). In this paper we introduce a multiple authentication scheme. That is, for *key* $:= (a, b, c, d) \in \mathbf{F}^4$ and $M \in \mathbf{F}$, let $\mathrm{auth}_4(M, \mathit{key}) := aM^3 + bM^2 + cM + d$. Note that the main advantage of the function $\mathrm{auth}_4()$ is that each authentication key $\mathit{key} = (a, b, c, d)$ can be used to authenticate three different messages M_0, M_1, and M_2 without revealing any information about any component of the authentication key. While for the function auth() each authentication key (a, b) can only be used to authenticate one message (that is, it is a kind of one-time pad) (see [15]), each authentication key (a, b, c, d) in our scheme can be used to authenticate three messages. Note that den Boer [4] used similar polynomials to construct one-time authentication schemes.

Theorem 2.1. *Let key* $:= (a, b, c, d)$ *be chosen uniformly from* $\mathbf{F}^4$, $M_i \in \mathbf{F}$, *and* $s_i := \mathrm{auth}_4(M_i, \mathit{key})$ *for* $i = 0, 1, 2$. *Then, for any* $\mathit{key}_0 := (a_0, b_0, c_0, d_0) \in \mathbf{F}^4$,

$$\begin{aligned}\Pr[a = a_0|\mathit{view}_0] &= \Pr[b = b_0|\mathit{view}_0] \\ &= \Pr[c = c_0|\mathit{view}_0] = \Pr[d = d_0|\mathit{view}_0] = \frac{1}{|\mathbf{F}|},\end{aligned}$$

where $\mathit{view}_0 := (M_0, s_0, M_1, s_1, M_2, s_2)$.

Proof. By the condition, one derives the following three equations with four unknowns:

$$\begin{aligned} M_0^3 a + M_0^2 b + cM_0 + d &= s_0, \\ M_1^3 a + M_1^2 b + cM_1 + d &= s_1, \\ M_2^3 a + M_2^2 b + cM_2 + d &= s_2. \end{aligned}$$

Since the coefficient matrix of the above equations is a Vandermonde matrix, no value of a can be ruled out. That is, every a is equally likely given the values $(M_0, s_0, M_1, s_1, M_2, s_2)$. (A similar argument applies for b, c, or d.) This completes the proof of the theorem. □

Following Franklin and Wright [7], we consider multicast as our only communication primitive. A message that is multicast by any node in a multicast neighbor network is received by all its neighbors with privacy (that is, nonneighbors learn nothing about what was sent) and authentication (that is, neighbors are guaranteed to receive the value that was multicast and to know which neighbor multicast it). We assume that all nodes in the multicast graph know the complete protocol specification and the complete structure of the multicast graph. In a message transmission protocol, the sender A starts with a message M^A drawn from a message space $\mathcal{M}$ with respect to a certain probability distribution. At the end of the protocol, the receiver B outputs a message M^B. We consider a synchronous system in which messages are sent via multicast in rounds. During each round of the protocol, each node receives any messages that were multicast by its neighbors at the end of the previous round, flips coins and performs local computations, and then possibly multicasts a message. We also assume that the message space $\mathcal{M}$ is a subset of a finite field $\mathbf{F}$.

We consider two kinds of adversaries. A passive adversary (or gossiper adversary) is an adversary who can only observe the traffic through t internal nodes. An active adversary (or Byzantine adversary) is an adversary with unlimited computational power who can control t internal nodes. That is, an active adversary will not only listen to the traffic through the controlled nodes, but also control the message sent by those controlled nodes. Both kinds of adversaries are assumed to know the complete protocol specification, message space, and the complete structure of the multicast graph. At the start of the protocol, the adversary chooses the t faulty nodes. A passive adversary can view the behavior (coin flips, computations, message received) of all the faulty nodes. An active adversary can view all the behavior of the faulty nodes and, in addition, control the message that they multicast. We allow for the stronger adversary. (An alternative interpretation is that t nodes are collaborating adversaries.)

For any execution of the protocol, let adv be the adversary's view of the entire protocol. We write $adv(M, r)$ to denote the adversary's view when $M^A = M$ and when the sequence of coin flips used by the adversary is r.

Definition 2.2 (see [7]).

1. A message transmission protocol is δ-*reliable* if, with probability at least $1 - \delta$, B terminates with $M^B = M^A$. The probability is over the choices of M^A and the coin flips of all nodes.

2. A message transmission protocol is ε-*private* if, for every two messages M_0, M_1 and every r, $\sum_c |\Pr[adv(M_0, r) = c] - \Pr[adv(M_1, r) = c]| \leq 2\varepsilon$. The probabilities are taken over the coin flips of the honest parties, and the sum is over all possible values of the adversary's view.
3. A message transmission protocol is *perfectly private* if it is 0-*private*.
4. A message transmission protocol is (ε, δ)-*secure* if it is ε-private and δ-reliable.
5. An (ε, δ)-secure message transmission protocol is *efficient* if its round complexity and bit complexity are polynomial in the size of the network, $\log(1/\varepsilon)$ (if $\varepsilon > 0$) and $\log(1/\delta)$ (if $\delta > 0$).

3. Background: Reliable Communication over Neighbor Networks

In this section we review Franklin and Wright's [7] protocols for reliable communication over multicast lines. The reader familiar with these protocols can skip this section. For two vertices A and B in a multicast graph $G(V, E)$, we say that A and B are connected by n *interiorly neighborhood-disjoint lines* if there are n lines $p_1, \ldots, p_n \subseteq V$ with the following properties:

- For each $1 \leq j \leq n$, the jth line p_j is a sequence of $m_j + 2$ nodes $A = X_{0,j}, X_{1,j}, \ldots, X_{m_j+1,j} = B$ where $X_{i,j}$ is a neighbor of $X_{i+1,j}$.
- For each i_1, i_2, j_1, and j_2 with $j_1 \neq j_2$, the only possible common neighbors of X_{i_1,j_1} and X_{i_2,j_2} are A and B.

Without loss of generality, in this section we assume that party A (the message transmitter) and party B (the message recipient) are connected by n interiorly neighborhood-disjoint lines, and we assume that $m_1 = m_2 = \cdots = m_n$.

Basic Propagation Protocol [7]. In this protocol, A tries to propagate a value s^A to B.

- In round 1, A multicasts s^A.
- In round ρ for $2 \leq \rho \leq m+1$, each $X_{\rho-1,j}$ $(1 \leq j \leq n)$ expects to receive a single element from $X_{\rho-2,j}$. Let $u_{\rho-1,j}$ be this value if a value was in fact received, or a publicly known default element otherwise. At the end of round ρ, $X_{\rho-1,j}$ multicasts $u_{\rho-1,j}$.
- In round $m + 2$, B receives a single element from each $X_{m,j}$, or substitutes the default element. Let s_j^B be the value received or substituted on line j.

From now on when a party substitutes the default element, we just say that the party substitutes.

Full Distribution Protocol [7]. In this protocol, each internal node $X_{i,j}$ tries to transmit an element $s_{i,j}$ to both A and B.

- In round 1, each $X_{i,j}$ $(1 \leq i \leq m, 1 \leq j \leq n)$ multicast $s_{i,j}$ to $X_{i-1,j}$ and $X_{i+1,j}$.

- In round ρ for $2 \leq \rho \leq m + 1$:

— For $1 \leq j \leq n$ and $\rho \leq i \leq m$, each $X_{i,j}$ expects to be the intended recipient of an element from $X_{i-1,j}$ (initiated by $X_{i-\rho+1,j}$). Let $u_{i,j}$ be the received value or a default value if none is received.

— For $1 \leq j \leq n$ and $1 \leq i \leq m - \rho + 1$, $X_{i,j}$ expects to be the intended recipient of an element from $X_{i+1,j}$ (initiated by $X_{i+\rho-1,j}$). Let $v_{i,j}$ be the received value or a default value if none if received.

— For $1 \leq j \leq n$, B expects to be the intended recipient on the jth line of a single element (initiated by $X_{m-\rho+2,j}$). Let $s^B_{m-\rho+2,j}$ be the received value or a default value if none is received.

— For $1 \leq j \leq n$, A expects to be the intended recipient on the jth line of a single element (initiated by $X_{\rho-1,j}$). Let $s^A_{\rho-1,j}$ be the received value or a default value if none is received.

— $X_{i,j}$ multicasts $u_{i,j}$ to $X_{i+1,j}$ if $\rho \leq i \leq m$, and $v_{i,j}$ to $X_{i-1,j}$ if $1 \leq i \leq m-\rho+1$.

Fact 3.1 [7]. *If there are no faults on the jth line, then $s^A_{i,j} = s^B_{i,j}$ for all $1 \leq i \leq m$. Further, if $X_{i,j}$ is the only fault on the jth line, then $s^A_{i,j} = s^B_{i,j}$.*

Reliable Transmission Protocol [7]. In this protocol, A reliably transmits a message M^A to B.

- The nodes on all the n lines execute an instance of the Full Distribution Protocol, which takes place during rounds 1 through $m + 1$. The element that $X_{i,j}$ initiates is $(a_{i,j}, b_{i,j})$ which is randomly chosen from $\mathbf{F}^2$. Let $(a^A_{i,j}, b^A_{i,j})$ and $(a^B_{i,j}, b^B_{i,j})$ be the values that A and B receive or substitute as the element initiated by $X_{i,j}$.
- The nodes on all the n lines execute an instance of the Basic Propagation Protocol from A to B, which takes place during rounds $m+2$ through $2m+3$. The element that A initiates is $\{(i, j, M^A, \text{auth}(M^A, a^A_{i,j}, b^A_{i,j})) : 1 \leq i \leq m, 1 \leq j \leq n\}$. In round $2m + 3$, B receives or substitutes $\{(i, j, M^B_{i,j,k}, u^B_{i,j,k}) : 1 \leq i \leq m, 1 \leq j \leq n\}$ on the kth line, $1 \leq k \leq n$.
- Let $r_k(M) := \{j : \exists i(M = M^B_{i,j,k}$ and $u^B_{i,j,k} = \text{auth}(M^B_{i,j,k}, a^B_{i,j}, b^B_{i,j}))\}$, that is, for any message M, $r_k(M)$ denotes the set of all line indices j such that M is "correctly" authenticated by some keys $(a^B_{i,j}, b^B_{i,j})$ according to the information B received on the kth line. B outputs M^B that maximizes $\max_k |r_k(M^B)|$.

Fact 3.2 [7]. *If $\delta > 0$, $n > t$, and $|\mathbf{F}| > mn^2/\delta$, then the Reliable Transmission Protocol is an efficient δ-reliable message transmission protocol.*

4. Reliable and Private Communication over Neighbor Networks

4.1. *Survey of Franklin–Wright's Results*

As in the previous section, we assume that party A (the message transmitter) and party B (the message recipient) are connected by n interiorly neighborhood-disjoint lines. Franklin and Wright [7] showed the following results regarding privacy in multicast networks:

1. If $n > t$, $\delta > 0$, and $\varepsilon > 0$, then there is an efficient (ε, δ)-secure message transmission protocol between A and B.
2. If $n > \lceil 3t/2 \rceil$ and $\delta > 0$, then there is an efficient $(0, \delta)$-secure message transmission protocol between A and B, that is, a δ-reliable and perfectly private message transmission protocol.
3. If $t < n \leq \lceil 3t/2 \rceil$ and $\delta > 0$, then there is an exponential bit complexity $(0, \delta)$-secure message transmission protocol between A and B.

4.2. *The Franklin–Wright Open Problem*

Franklin and Wright left open the question whether it is possible to achieve perfect privacy efficiently when $t < n \leq \lceil 3t/2 \rceil$. That is, does there exist a polynomial time $(0, \delta)$-secure message transmission protocol between A and B when $t < n \leq \lceil 3t/2 \rceil$? We give an affirmative answer to this question.

4.3. *The Solution*

Intuitively, our protocol proceeds as follows. First, using the Full Distribution Protocol from the preceding section, each internal node $X_{i,j}$ transmits a random authentication key $\mathit{key}_{i,j} = (a_{i,j}, b_{i,j}, c_{i,j}, d_{i,j}) \in_R \mathbf{F}^4$ to both A and B. Secondly, using the Basic Propagation Protocol, B transmits to A a random $r_{i,j} \in_R \mathbf{F}$ authenticated by the key $\mathit{key}_{i,j}$ for each $1 \leq i \leq m$, $1 \leq j \leq n$. Thirdly, for each $1 \leq j \leq n$, A decides whether A and B agree on at least one authentication key on the jth line. Informally, let

$$\begin{aligned} K^A := \{(i_j, j) : &A \text{ believes that } \mathit{key}^A_{i_j,j} \text{ is the first key} \\ &\text{agreed upon by} A \text{ and} B \text{ on the} j\text{th line}\}. \end{aligned}$$

The formal definition of K^A is given in the following protocol. Fourthly, A encrypts the message M^A using the sum of the pads $a^A_{i_j,j}$ $((i_j, j) \in K^A)$ and, using the Reliable Transmission Protocol, transmits to B the set K^A and the ciphertext. Lastly, B decrypts the message.

Perfectly Private Transmission Protocol.

- The nodes on all the n lines execute an instance of the Full Distribution Protocol, which takes place during rounds 1 through $m + 1$. The element that $X_{i,j}$ initiates is $\mathit{key}_{i,j} = (a_{i,j}, b_{i,j}, c_{i,j}, d_{i,j})$ which is randomly chosen from $\mathbf{F}^4$. Let $\mathit{key}^A_{i,j} := (a^A_{i,j}, b^A_{i,j}, c^A_{i,j}, d^A_{i,j})$ and $\mathit{key}^B_{i,j} := (a^B_{i,j}, b^B_{i,j}, c^B_{i,j}, d^B_{i,j})$ be the values that A and B receive or substitute as the element initiated by $X_{i,j}$.
- For each i, j, B chooses $r^B_{i,j} \in_R \mathbf{F}$. The nodes on all the n lines execute an instance of the Basic Propagation Protocol from B to A, which takes place during rounds $m + 2$ through $2m + 3$. The element that B initiates is

  $$\langle s^B, \langle \mathrm{auth}_4(s^B, \mathit{key}^B_{i,j}) : 1 \leq i \leq m, 1 \leq j \leq n \rangle \rangle,$$

  where $s^B := \langle \langle r^B_{i,j}, \mathrm{auth}_4(r^B_{i,j}, \mathit{key}^B_{i,j}) \rangle : 1 \leq i \leq m, 1 \leq j \leq n \rangle$ and $\langle \cdots \rangle$ denotes the ordered set of its elements (without loss of generality, we assume that we can uniquely and efficiently recover its elements from the ordered set $\langle\cdots\rangle$). In round $2m+3$, A receives or substitutes $\langle s_k^A, \langle u_{i,j,k}^A : 1 \le i \le m, 1 \le j \le n\rangle\rangle$ on the kth line for $1 \le k \le n$. For each triple (i, j, k), let $v_{i,j,k}^A$ be defined in such a way that $s_k^A := \langle\langle r_{i,j,k}^A, v_{i,j,k}^A\rangle : 1 \le i \le m, 1 \le j \le n\rangle$.

- Let $L_k := \{j : \exists i(u_{i,j,k}^A = \mathrm{auth}_4(s_k^A, key_{i,j}^A))\}$, that is, L_k denotes the set of all line indices j such that s_k^A is correctly authenticated by some keys $key_{i,j}^A$ according to the information A received on the kth line, and let k_0 be the line that maximizes $|L_{k_0}| = \max_k |L_k|$. Let $K^A := \{(i_j, j) : \forall(0 < i < i_j)(v_{i,j,k_0}^A \neq \mathrm{auth}_4(r_{i,j,k_0}^A, key_{i,j}^A))$ and $v_{i_j,j,k_0}^A = \mathrm{auth}_4(r_{i_j,j,k_0}^A, key_{i,j}^A)\}$. If $|K^A| \ge |L_{k_0}|$, then go to the next step, otherwise, A terminates the protocol with failure.
- A computes $z^A = M^A + \sum_{(i_j,j)\in K^A} a_{i_j,j}^A$.
- In rounds $2m+4$ through $4m+7$, the nodes on all the n lines execute an instance of the Reliable Transmission Protocol from A to B. The element that A initiates is $\langle z^A, \langle K^A\rangle\rangle$. Let $\langle z^B, \langle K^B\rangle\rangle$ be the value received by B as the output of the Reliable Transmission Protocol.
- B computes $M^B := z^B - \sum_{(i_j,j)\in K^B} a_{i_j,j}^B$.

The Perfectly Private Transmission Protocol provides efficient $(0, \delta)$-secure message transmission provided that the field $\mathbf{F}$ used by $\mathrm{auth}_4()$ satisfies $|\mathbf{F}| \ge mn(2n+3)/\delta$. Since reliable communication is not possible when $t \ge n$, this protocol provides matching upper and lower bounds for perfect privacy and probabilistic reliability.

Theorem 4.1. *If $\delta > 0$, $n > t$, and $|\mathbf{F}| > mn(2n+3)/\delta$, then the Perfectly Private Transmission Protocol is an efficient $(0, \delta)$-secure message transmission protocol.*

Proof. To see that the Perfectly Private Transmission Protocol is 0-private, let w_0 denote the number of lines with no faults, w_1 the number with exactly one fault, and w_+ the number with two or more faults. Then $n = w_0 + w_1 + w_+$ and $t \ge w_1 + 2w_+$. Since $n > t$, it follows that $w_0 > w_+$. By Fact 3.1, no matter whether the transmission from B to A during rounds $m+2$ to $2m+3$ succeeds or not, $|K^A| \ge |L_{k_0}| \ge w_0 + w_1 > w_+ + w_1$. Whence there is a $(i_l, l) \in K^A$ such that the lth line is a nonfaulty line, and $key_{i_l,l}^A = key_{i_l,l}^B$. By Theorem 2.1, the adversary gets no information about $a_{i_l,l}^A$ given the view adv_{M^A}, where adv_{M^A} consists of the following information:

1. $\langle\langle r_{i,j}^B, \mathrm{auth}_4(r_{i,j}^B, key_j^B)\rangle : 1 \le i \le m, 1 \le j \le n\rangle$;
2. $\langle s^B, \langle \mathrm{auth}_4(s^B, key_{i,j}^B) : 1 \le i \le m, 1 \le j \le n\rangle\rangle$;
3. at most one randomly guessed (by the adversary) correct authenticator of some random message.

It should be noted that the above item 3 in the adversary's view adv_{M^A} is important for the following reasons: with nonzero probability the first transmission from B to A may fail (i.e., in rounds $m+2$ through $2m+3$). That is, the adversary may create a bogus $(r_{i,j}^B)'$ (which is different from $r_{i,j}^B$) and guess the value $\mathrm{auth}_4((r_{i,j}^B)', key_{i,j}^B)$ correctly. Then at the end of round $2m+3$, A may choose a wrong K^A. That is, there may be an item $(i_{j'}, j') \in K^A$ such that $key_{i_{j'},j'}^A$ is not the first key agreed upon by A and B on the j'th line. It is easy for the adversary to decide whether such kind of an item exists in

K^A. When such an item exists, the adversary knows that he may have guessed a correct authenticator of the message $(r^B_{i,j})'$. Since

$$z^A = M^A + a^A_{i_l,l} + \sum_{\substack{(i_j,j)\in K^A \\ j\neq l}} a^A_{i_j,j},$$

one derives that every M^A is equally likely given adv_{M^A}. Since this is the only relevant information about M^A in *adv*, one derives that $\Pr[adv(M_0, r) = c] = \Pr[adv(M_1, r) = c]$ for every pair of messages M_0 and M_1, adversary's coin flips r, and the possible view c. It follows that

$$\sum_c |\Pr[adv(M_0, r) = c] - \Pr[adv(M_1, r) = c]| = 0.$$

In the following we prove reliability. Let

$$\begin{aligned} K^{AB} := \{(i_j, j) : \quad & \exists i\,(key^A_{i,j} = key^B_{i,j}) \text{ and} \\ & \forall (0 < i < i_j)(key^A_{i,j} \neq key^B_{i,j})\}, \end{aligned}$$

$\text{KEY}^A := \sum_{(i,j)\in K^A} a^A_{i,j}$, and $\text{KEY}^{AB} := \sum_{(i,j)\in K^{AB}} a^A_{i,j}$. Note that, by a simple argument as in the proof of Fact 3.2, it is straightforward that the first transmission from B to A (i.e., in rounds $m+2$ through $2m+3$) succeeds with probability at least $1 - mn^2/|\mathbf{F}|$. Let FRT denote the event that the first transmission from B to A (i.e., in rounds $m+2$ through $2m+3$) succeeds. In the following we compute the probability that $\text{KEY}^A = \text{KEY}^{AB}$ given FRT. We first define two notations. For any (i, j), let

$$T^A(i, j) := \begin{cases} 1 & \text{if} \quad v^A_{i,j,k_0} = \text{auth}_4(r^A_{i,j,k_0}, key^A_{i,j}), \\ 0 & \text{otherwise;} \end{cases}$$

and

$$T^{AB}(i, j) := \begin{cases} 1 & \text{if} \quad key^A_{i,j} = key^B_{i,j}, \\ 0 & \text{otherwise.} \end{cases}$$

It is straightforward that for any (i, j), one derives

$$\Pr[T^{AB}(i, j) = T^A(i, j) | \text{ FRT}, T^{AB}(i, j) = 1] = 1. \tag{1}$$

For any (i, j) such that $T^{AB}(i, j) = 0$ and $T^A(i, j) = 1$, one derives $v^A_{i,j,k_0} = \text{auth}_4(r^A_{i,j,k_0}, key^A_{i,j})$. Since FRT implies that $v^A_{i,j,k_0} = \text{auth}_4(r^B_{i,j}, key^B_{i,j})$ and $r^B_{i,j} = r^A_{i,j,k_0}$, one derives

$$a^B_{i,j}(r^B_{i,j})^3 + b^B_{i,j}(r^B_{i,j})^2 + c^B_{i,j}r^B_{i,j} + d^B_{i,j} = a^A_{i,j}(r^A_{i,j,k_0})^3 + b^A_{i,j}(r^A_{i,j,k_0})^2 + c^A_{i,j}r^A_{i,j,k_0} + d^A_{i,j},$$

which implies that $r^B_{i,j}$ is a solution of the equation

$$(a^B_{i,j} - a^A_{i,j})(r^B_{i,j})^3 + (b^B_{i,j} - b^A_{i,j})(r^B_{i,j})^2 + (c^B_{i,j} - c^A_{i,j})r^B_{i,j} + (d^B_{i,j} - d^A_{i,j}) = 0. \tag{2}$$

Since $key^A_{i,j}$ and $key^B_{i,j}$ are fixed before the random choice of $r^B_{i,j}$, and (2) has at most three solutions, we have the following relation:

$$\begin{aligned}&\Pr[T^{AB}(i,j) = T^A(i,j)|\text{FRT}, T^{AB}(i,j) = 0, r^B_{i,j}]\\ &\quad = \begin{cases}0 & \text{if} \quad r^B_{i,j} \text{ is a solution of (2),}\\ 1 & \text{otherwise.}\end{cases}\end{aligned}$$

This implies that

$$\begin{aligned}&\Pr[T^{AB}(i,j) = T^A(i,j)|\text{FRT}, T^{AB}(i,j) = 0]\\ &\quad = \sum_{r^B_{i,j}\in\mathbf{F}} \Pr[T^{AB}(i,j) = T^A(i,j)|\text{FRT}, T^{AB}(i,j) = 0, r^B_{i,j}] \cdot \Pr[r^B_{i,j}]\\ &\quad \geq 1 - \frac{3}{|\mathbf{F}|}. \end{aligned} \tag{3}$$

Combining (1) and (3), for any (i, j) one derives

$$\Pr[T^{AB}(i,j) = T^A(i,j)|\,\text{FRT}] \geq 1 - \frac{3}{|\mathbf{F}|}.$$

Hence we have

$$\begin{aligned}\Pr\big[\text{KEY}^A = \text{KEY}^{AB}\,|\,\text{FRT}\big] &\geq \prod_{\substack{1\leq i\leq m\\ 1\leq j\leq n}} \Pr\big[T^A(i,j) = T^{AB}(i,j)|\,\text{FRT}\big]\\ &\geq \prod_{\substack{1\leq i\leq m\\ 1\leq j\leq n}} \left(1 - \frac{3}{|\mathbf{F}|}\right)\\ &\geq 1 - \frac{3mn}{|\mathbf{F}|}.\end{aligned}$$

Note that the last inequality is obtained by the Bernoulli inequality, that is, $(1+x)^{mn} \geq 1+mnx$ for any $x \geq -1$. Let SRT denote the event that the second reliable transmission from A to B (i.e., in rounds $2m+4$ through $4m+7$) succeeds. Then

$$\begin{aligned}\Pr[M^A = M^B] &\geq \Pr[\text{FRT} \wedge \text{SRT} \wedge (\text{KEY}^A = \text{KEY}^{AB})]\\ &\geq \Pr[(\text{KEY}^A = \text{KEY}^{AB})|\,\text{FRT, SRT}] \cdot \Pr[\text{FRT, SRT}]\\ &= \Pr\big[\text{KEY}^A = \text{KEY}^{AB}\,|\,\text{FRT}\big] \cdot \Pr[\text{FRT}] \cdot \Pr[\text{SRT}]\\ &\geq \left(1 - \frac{3mn}{|\mathbf{F}|}\right)\left(1 - \frac{mn^2}{|\mathbf{F}|}\right)^2\\ &\geq 1 - \frac{mn(2n+3)}{|\mathbf{F}|}.\end{aligned}$$

Note that we have used the following properties in the above computation:

1. FRT and SRT are independent;
2. SRT and the event $(\text{KEY}^A = \text{KEY}^{AB})$ are independent.

Since $|\mathbf{F}| > mn(2n+3)/\delta$, it follows that $\Pr[M^B = M^A] > 1 - \delta$. □

5. Weak Connectivity

In the more general setting of multicast graph, there is a channel from each node to its neighbor nodes. We say that two nodes A and B of a multicast graph are *strongly t-connected* (which was implicitly introduced by Franklin and Wright [7]) if there are t interiorly neighborhood-disjoint paths connecting A and B. Franklin and Wright [7] have observed that the multicast lines protocol can be simulated on any strongly $(t+1)$-connected multicast graph. That is, if A and B are strongly $(t+1)$-connected, then our result in the previous section shows that $(0, \delta)$-secure message transmission between A and B is possible. In the following we show that this condition is not necessary.

Franklin and Yung [8] define that two nodes A and B in a multicast graph $G(V, E)$ are *weakly t-connected* if for any set $V_1 \subseteq V \backslash \{A, B\}$ with $|V_1| < t$, the removal of $neighbor(V_1)$ and all incident edges from $G(V, E)$ does not disconnect A and B, where $neighbor(V_1) = V_1 \cup \{v \in V : \exists u \in V_1(u, v) \in E\} \backslash \{A, B\}$. Franklin and Yung [8] show that it is co**NP** hard to decide whether a given graph is weakly t-connected.

Let A and B be two nodes on a multicast graph $G(V, E)$ and $t < n$. We say that A and B are *weakly (n, t)-connected* if there are n vertex disjoint paths $p_1, \ldots, p_n$ between A and B and, for any vertex set $T \subseteq (V \backslash \{A, B\})$ with $|T| \leq t$, there exists an i $(1 \leq i \leq n)$ such that all vertices of p_i have no neighbor in T. Obviously, if two vertices are weakly (n, t)-connected then they are weakly $(t+1)$-connected.

Theorem 5.1. *If A and B are* weakly (n, t)-connected *for some $t < n$, then the Perfectly Private Transmission Protocol in the previous section is an efficient $(0, \delta)$-secure message transmission between A and B.*

Proof. It follows straightforwardly from the proof of Theorem 4.1. □

Franklin and Yung [8] show that, in the context of a t-passive adversary, weak $(t+1)$-connectivity is necessary and sufficient for achieving private communications. Theorem 5.1 provides a sufficient condition for achieving perfect privacy and probabilistic reliability against a t-active adversary in a general multicast graph. It is an open question whether the condition in Theorem 5.1 is also necessary.

It is easily observed that strong $(t+1)$-connectivity implies weak $(t+1, t)$-connectivity. The following example shows that (n, t)-weak connectivity does not imply strong $(t+1)$-connectivity.

Example 5.2. Let $G(V, E)$ be the graph defined by $V := \{A, B\} \cup \{v_{i,j} : i, j = 1, 2, 3\}$ and $E := \{(A, v_{1,j}) : j = 1, 2, 3\} \cup \{(v_{i,j}, v_{i+1,j}) : i = 1, 2;\ j = 1, 2, 3\} \cup \{(v_{3,j}, B) : j = 1, 2, 3\} \cup \{(v_{1,1}, v_{1,2}), (v_{2,2}, v_{2,3}), (v_{3,3}, v_{3,1})\}$. The graph G is graphically displayed in Fig. 1. Then it is straightforward to show that A and B are weakly $(3, 1)$-connected but not strongly 2-connected in G.

Theorem 5.1 shows that, for at most one malicious node, efficient $(0, \delta)$-secure message transmission between A and B is possible in the multicast graph defined in Example 5.2. Note that this multicast graph is only strongly 1-connected, and so Franklin–Wright's results have no bearing on this example.

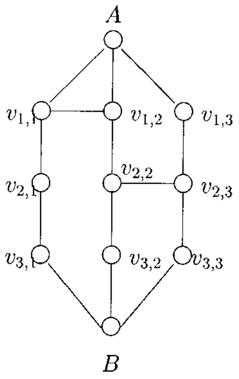

Fig. 1. Figure for Example 5.2.

Similarly, for any $n > 2$ the following example gives a graph G and two vertices A and B such that A and B are weakly $(n, 1)$-connected but not weakly 3-connected.

Example 5.3. Let $G(V, E)$ be the graph defined by $V = \{A, B\} \cup \{v_{i,j} : i = 1, 2; j = 1, \ldots, n\}$ and $E = \{(A, v_{1,j}) : j = 1, \ldots, n\} \cup \{(v_{1,j}, v_{2,j}) : j = 1, \ldots, n\} \cup \{(v_{2,j}, B) : j = 1, \ldots, n\} \cup \{(v_{2,1}, v_{2,j}) : j = 2, \ldots, \lfloor n/2 \rfloor\} \cup \{(v_{2,\lfloor n/2 \rfloor+1}, v_{2,j}) : j = \lfloor n/2 \rfloor + 2, \ldots, n\}$. The graph G is graphically displayed in Fig. 2. Then it is straightforward to show that A and B are weakly $(n, 1)$-connected but not weakly 3-connected in G.

Then Theorem 5.1 shows that, for at most one malicious node, efficient $(0, \delta)$-secure message transmission between A and B is possible in the graph G defined in Example 5.3. The result by Franklin and Yung [8] shows that secure message transmission between A and B is impossible in this graph when there are two malicious nodes. However, if $n > 2t + 1$ and we use nonmulticast channels, then secure message transmission is possible between A and B against t malicious nodes (see, e.g., [6]). It follows that in certain cases multicast *helps* adversaries "more," which contrasts with Franklin and Wright's result [7] that in certain cases multicast *hurts* adversaries "more."

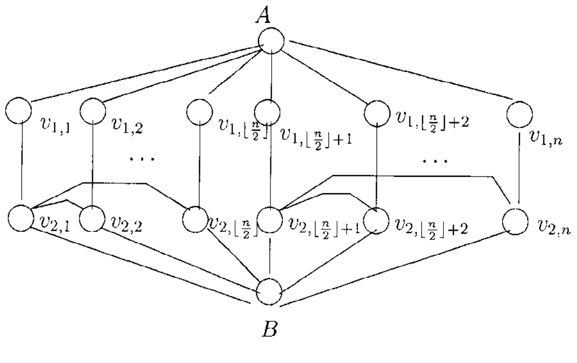
A
$v_{1,1}$
$v_{1,2}$
...
$v_{1,\lfloor \frac{n}{2} \rfloor}$
$v_{1,\lfloor \frac{n}{2} \rfloor+1}$
$v_{1,\lfloor \frac{n}{2} \rfloor+2}$
...
$v_{1,n}$
$v_{2,1}$
$v_{2,2}$
$v_{2,\lfloor \frac{n}{2} \rfloor}$
$v_{2,\lfloor \frac{n}{2} \rfloor+1}$
$v_{2,\lfloor \frac{n}{2} \rfloor+2}$
$v_{2,n}$
B

Fig. 2. Figure for Example 5.3.

We close our paper by showing that it is **NP**-hard to decide whether a given multicast graph is strongly t-connected. Given a graph $G(V, E)$, We say that a subset V_1 of V is *neighborhood independent* if any two nodes in V_1 have no common neighbor in G. We first prove the following lemma.

Lemma 5.4. *Given a graph $G(V, E)$ and a number t, it is* **NP***-complete to decide whether there exists a neighborhood independent set $V_1 \subseteq V$ of size t.*

Proof. It is clear that the specified problem is in **NP**. Whence it suffices to reduce the following **NP**-complete problem IS (Independent Set) to our problem. The independent set problem is:

Instance: A graph $GI(V_G, E_G)$ and a number t.
Question: Does there exist a node set $V_1 \subseteq V_G$ of size t such that any two nodes in V_1 are not connected by an edge in E_G?

The input $GI(V_G, E_G)$, to IS, consists of a set of vertices $V_G = \{v_1, \dots, v_n\}$ and a set of edges E_G. In the following we construct a graph $f(GI) := G(V, E)$ such that there is an independent set of size t in GI if and only if there is a neighborhood independent set of size t in G.

Let $V := V_G \cup V'$ where $V' = \{v_{i,j} : (v_i, v_j) \in E_G, i < j\}$, and $E := \{(v_i, v_{i,j}), (v_{i,j}, v_j) : v_{i,j} \in V'\} \cup \{(v_{i,j}, v_{i',j'}) : v_{i,j}, v_{i',j'} \in V'\}$. It is straightforward to check that, for any neighborhood independent set $V_1 \subseteq V$, if $V_1 \cap V' \neq \emptyset$ then $|V_1| = 1$. It is also clear that for any two vertex $u, v \in V_G$, u and v have no common neighbor in $f(GI)$ if and only if $(u, v) \notin E_G$. Hence there is a neighborhood independent set of size t in G if and only if there is an independent set of size t in GI. □

Theorem 5.5. *It is* **NP***-complete to decide whether a given multicast graph is strongly t-connected.*

Proof. It is clear that the specified problem is in **NP**. Whence it suffices to reduce the **NP**-complete problem NIS (Neighborhood Independent Set) in Lemma 5.4 to our problem.

The input $G(V_G, E_G)$, to NIS, consists of a set of vertices $V_G = \{v_1, \dots, v_n\}$ and a set of edges E_G. In the following we construct a multicast graph $f(G) = MG(V, E)$ and two nodes $A, B \in V$ such that there is a neighborhood independent set of size t in G if and only if A and B are strongly t-connected.

Let $V := \{A, B\} \cup V_G$ and $E := E_G \cup \{(A, v), (v, B) : v \in V_G\}$. It is clear that two paths P_1 and P_2 connecting A and B which go through v_i and v_j respectively are node disjoint and have no common neighbor (except A and B) if and only if v_i and v_j have no common neighbor in $G(V_G, E_G)$. Hence there is a neighborhood independent set of size t in G if and only if A and B are strongly t-connected in $f(G)$. □

Similarly, we can show that the corresponding problem for weak (n, t)-connectivity is co**NP**-hard.

Theorem 5.6. *It is co***NP***-hard to decide whether a given multicast graph is weakly* (n, t)*-connected.*

Proof. It suffices to reduce the following co**NP**-complete problem coVC (co-Vertex Cover) to our problem. The co-VC problem is:

Instance: A graph $GC(V_G, E_G)$ and a number t.
Question: Does there not exist a node set $V_1 \subseteq V_G$ of size t such that for each edge $(u, v) \in E_G$, at least one of u and v belongs to V_1?

The input $GC(V_G, E_G)$, to coVC, consists of a set of vertices $V_G = \{v_1, \dots, v_m\}$ and a set of edges $E_G = \{e_1, \dots, e_n\}$. In the following we construct a graph $f(GC) := G(V, E)$ such that there does not exist a vertex cover of size t in GC if and only if $f(GC)$ is weakly (n, t)-connected.

Let $V := V_G \cup V' \cup \{A, B\}$ where $V' = \{v_{i,j} : (v_i, v_j) \in E_G, i < j\}$, and $E := \{(A, v_{i,j}), (v_{i,j}, B) : v_{i,j} \in V'\} \cup \{(v_i, v_{i,j}), (v_{i,j}, v_j) : v_{i,j} \in V'\}$. For any t-size subset $V_1 \subseteq V$, let $V_2 = (V_1 \cap V_G) \cup \{v_i : (v_i, v_j) \in V_1\}$. Then it is straightforward to check that the following condition holds:

- There exists a pair (i, j) such that $v_{i,j}$ has no neighbor in V_1 if and only if V_2 is not a vertex cover of $GC(V_G, E_G)$.

Hence $f(GC)$ is weakly (n, t)-connected (witnesses by the vertex disjoint n paths: $A \to v_{i,j} \to B$ $(v_{i,j} \in V')$) if and only if GC does not have a vertex cover of size t. □

Indeed, it is straightforward to show that the above problem belongs to Σ_2^p (that is, the second level of the polynomial time hierarchy). It remains open whether this problem is co**NP**-complete, or Σ_2^p-complete, or neither.

Acknowledgments

The authors thank Matt Franklin (Xerox) and Rebecca Wright (AT&T) for informing us that Donald Beaver has found a different method to address the Franklin–Wright open problem. The authors would also like to thank two anonymous referees for their detailed comments on the presentation of this paper.

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, ACM Press, New York, 1988.

[2] M. Burmester, Y. Desmedt, and G. Kabatianski. Trust and security: a new look at the Byzantine generals problem. DIMACS Series in Discrete Mathematics and Theoretical Computer Science 38, pages 75–83, American Mathematical Society, Providence, RI, 1998.

[3] D. Chaum, C. Crepeau, and I. Damgard. Multiparty unconditional secure protocols. In: *Proc. ACM STOC* '88, pages 11–19, ACM Press, New York, 1988.

[4] B. den Boer. A simple and key-economical unconditional authentication scheme. *J. Comput. Security,* **2**:65–71, 1993.

[5] D. Dolev. The Byzantine generals strike again. *J. Algorithms*, **3**:14–30, 1982.

[6] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *J. Assoc. Comput. Mach.*, **40**(1):17–47, 1993.

[7] M. Franklin and R. Wright. Secure communication in minimal connectivity models. *J. Cryptology*, **13**(1):9–30, 2000.

[8] M. Franklin and M. Yung. Secure hypergraphs: privacy from partial broadcast. In: *Proc. ACM STOC* '95, pages 36–44, ACM Press, New York, 1995.

[9] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of* **NP**-*Completeness*. Freeman, San Francisco, CA, 1979.

[10] E. Gilbert, F. MacWilliams, and N. Sloane. Codes which detect deception. *Bell System Tech. J.*, **53**(3):405–424, 1974.

[11] O. Goldreich, S. Goldwasser, and N. Linial. Fault-tolerant computation in the full information model. *SIAM J. Comput.*, **27**(2):506–544, 1998.

[12] V. Hadzilacos. Issues of Fault Tolerance in Concurrent Computations. Ph.D. thesis, Harvard University, Cambridge, MA, 1984.

[13] T. Rabin. Robust sharing of secrets when the dealer is honest or faulty. *J. Assoc. Comput. Mach.*, **41**(6):1089–1109, 1994.

[14] T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In: *Proc. ACM STOC* '89, pages 73–85, ACM Press, New York, 1989.

[15] G. J. Simmons. A survey of information authentication. In: *Contemporary Cryptology, The Science of Information Integrity*, pages 379–419. IEEE Press, New York, 1992.

[16] M. Yannakakis. Node- and edge-deletion **NP**-complete problems. In: *Proc. ACM STOC* '78, pages 253–264, ACM Press, New York, 1978.