

Approximation Hardness And Secure Communication In Broadcast Channels

Yvo Desmedt^{1,2*}, Yongge Wang^{3**}

¹ Computer Science, Florida State University, Tallahassee
Florida FL 32306-4530, USA, desmedt@cs.fsu.edu

² Dept. of Mathematics, Royal Holloway, University of London, UK

³ Center for Applied Cryptographic Research, Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada,
ygwang@cacr.math.uwaterloo.ca

Abstract. Problems of secure communication and computation have been studied extensively in network models. Goldreich, Goldwasser, and Linial, Franklin and Yung, and Franklin and Wright have initiated the study of secure communication and secure computation in multi-recipient (broadcast) models. A “broadcast channel” (such as Ethernet) enables one processor to send the same message—simultaneously and privately—to a fixed subset of processors. Franklin and Wright, and Wang and Desmedt have shown that if there are at most k malicious (Byzantine style) processors, then there is an efficient protocol for achieving probabilistically reliable and perfectly private communication in a strongly n -connected network where $n \geq k + 1$. While these results are unconditional, we will consider these problems in the scenario of conditional reliability, and then improve the bounds. In this paper, using the results for hardness of approximation and optimization problems, we will design communication protocols (with broadcast channels) which could defeat more faults than possible with the state of the art. Specifically, assuming certain approximation hardness result, we will construct strongly n -connected graphs which could defeat a k -active adversary (whose computation power is polynomially bounded) for $k = cn$, where $c > 1$ is any given constant. This result improves a great deal on the results of Franklin and Wright, and Wang and Desmedt.

* Research partly supported by DARPA F30602-97-1-0205. However the views and conclusions contained in this paper are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advance Research Projects Agency (DARPA), the Air Force, of the US Government.

This research was done while the author was at the University of Wisconsin – Milwaukee.

** Most of the research was done when the author was a post docs of Professor Desmedt at the University of Wisconsin – Milwaukee.

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 have been studied extensively (see, e.g., [2, 4, 5, 10]). For example, Dolev [4] and Dolev, Dwork, Waarts, and Yung [5] showed that, in the case of k Byzantine faults, reliable communication is achievable only if the system's network is $2k + 1$ connected. Hadzilacos [10] has shown that even in the absence of malicious failures connectivity $k + 1$ is required to achieve reliable communication in the presence of k faulty participants.

Goldreich, Goldwasser, and Linial [9], Franklin and Yung [7], and Franklin and Wright [6] have initiated the study of secure communication and secure computation in *multi-recipient (broadcast)* models. A “broadcast channel” (such as Ethernet) enables one participant to send the same message—simultaneously and privately—to a fixed subset of participants. Franklin and Yung [7] have given a necessary and sufficient condition for individuals to exchange private messages in broadcast 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 [6]. Note that Goldreich, Goldwasser, and Linial [9] have also studied the fault-tolerant computation in the public broadcast model in the presence of active Byzantine adversaries.

There are many examples of broadcast 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 broadcast to the subset of participants that is able to decrypt it.

We will abstract away the concrete network structures and consider multicast graphs. Specifically, a multicast graph is just a graph $G(V, E)$. A vertex $A \in V$ is called a neighbor of another vertex $B \in V$ if 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 neighborhood learn nothing about the content of the message. The neighborhood networks have been studied by Franklin and Yung in [7]. They have also studied the more general notion of hypergraphs, which we do not need.

As Franklin and Wright [6] have pointed out, unlike the simple channel model, it is not possible to directly apply protocols 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 simple channel to broadcast 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.

It was shown [6] that if the sender and the receiver are strongly n -connected (that is, there are n paths with disjoint neighborhoods) and the malicious adver-

sary can destroy at most k processors, then the condition $n > k$ is necessary and sufficient for achieving efficient probabilisticly reliable and probabilisticly private communication. They also showed that there is an efficient protocol to achieve probabilisticly reliable and perfectly private communication when $n > \lceil 3k/2 \rceil$. Recently, Wang and Desmedt [13] have shown that, indeed, the condition $n > k$ is necessary and sufficient for achieving efficient probabilisticly reliable and perfectly private communication in broadcast channels.

Definition 1. *Let A and B be two vertices on a multicast graph $G(V, E)$. We say that A and B are strongly n -connected if there are n neighborhood disjoint paths p_1, \dots, p_n between A and B , that is, for any $i \neq j (\leq n)$, p_i and p_j have no common neighbor (except A and B). In other words, for any vertex $v \in V \setminus \{A, B\}$, if there is a vertex u_1 on p_i such that $(v, u_1) \in E$, then there is no u_2 on p_j such that $(v, u_2) \in E$.*

However, all these results are concerned with malicious adversaries with unlimited computational power. In this paper, we will consider the situation when the adversary's computational power is polynomial time bounded. Specifically, assuming certain approximation hardness result, we will construct strongly n -connected multicast graphs which could defeat a k -active adversary (whose computation power is polynomial time bounded) for $k = cn$, where $c > 1$ is any given constant. This result improves a great deal on the results of Franklin and Wright [6] (which are for unconditional reliability). To achieve this improvement we use some of the hardness results of Burmester, Desmedt, and Wang in [3].

The idea underlying our construction is that we will design strongly n -connected communication graphs in such a way that it is hard for the adversary to find the neighborhood disjoint n paths which is a witness to the strong n -connectivity. Hence the adversary does not know which processors to block (or control).

There have been many results (see, e.g., [1, 12] for a survey) for hardness of approximating an **NP**-hard optimization problem within a factor c from "below". For example, it is hard to compute an independent set¹ V' of a graph $G(V, E)$ with the property that $|V'| \geq \frac{n}{c}$ for some given factor c , where n is the size of the maximum independent set of G . But for our problem, we are more concerned with approximating an **NP**-hard optimization problem from "above". For example, given a graph $G(V, E)$, how hard is it to compute a vertex set V' of G with $|V'| \leq cn$ such that V' contains an optimal independent set of G , where n is the size of the optimal independent set of G ? Burmester, Desmedt, and Wang have shown that this kind of approximation problem is also **NP**-hard. We will use this result to design strongly n -connected multicast graphs which is secure against an active adversary who can control cn vertices where $c > 1$ is some constant.

The organization of this paper is as follows. We first present in Section 2 some graph theoretic result which we will need in this paper. Section 3 surveys

¹ An independent set in a graph $G(V, E)$ is a subset V' of V such that no two vertices in V' are joined by an edge in E .

the model for commutation in broadcast channels. In Section 4 we demonstrate how to use strongly n -connected graphs with trapdoors to achieve reliable and private communication against active adversaries. In Section 5 we outline an approach to build strongly n -connected graphs with trapdoors. We conclude in Section 6 with remarks towards theoretical improvements and we present some open problems.

2 Optimization and approximation

In this section we survey and introduce some graph theoretic results which will be used in later sections.

Definition 2. *The independent set problem is:*

Instance: A graph $G(V, E)$ and a number n .

Question: Does there exist a vertex set $V_1 \subseteq V$ of size n such that any two nodes in V_1 are not connected by an edge in E ?

Definition 3. *Given a graph $G(V, E)$, a vertex subset $V' \subseteq V$ is called neighborhood independent if for any $u, v \in V'$ there is no $w \in V$ such that both (u, w) and (v, w) are edges in E .*

Definition 4. *A vertex v in a graph $G(V, E)$ is isolated if there is no edge adjacent to v , i.e., for all $w \in V$, $(v, w) \notin E$.*

Theorem 1. *Given a graph $G(V, E)$ and a number n , it is **NP**-complete to decide whether there exists a neighborhood independent set $V_1 \subseteq V$ of size n .*

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

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

Let $V_G = V \cup V'$ where $V' = \{v_{i,j} : (v_i, v_j) \in E, i < j\} \cup \{v_{i,i} : v_i \text{ is an isolated vertex}\}$ and $E_G = \{(v_i, v_{i,j}), (v_{i,j}, v_j) : v_{i,j} \in V', i \leq j\} \cup \{(v_{i,j}, v_{i',j'}) : v_{i,j}, v_{i',j'} \in V', i \leq j, i' \leq j'\}$. It is straightforward to check that, for any neighborhood independent set $V_1 \subseteq V_G$, if $V_1 \cap V' \neq \emptyset$ then $|V_1| = 1$. It is also clear that for any two vertex $u, v \in V$, u and v have no common neighbor in $f(G)$ if and only if $(u, v) \notin E$. Hence there is a neighborhood independent set of size n in GNI if and only if there is an independent set of size n in G . \square

The following results follow directly from the corresponding results for independent sets in Burmester, Desmedt, and Wang [3].

Theorem 2. ([3]) *There is a constant $\varepsilon > 0$ such that it is **NP**-hard to compute a vertex set $V' \subseteq V$ of a graph $G(V, E)$, with the following properties:*

1. $|V'| \leq nm^\varepsilon$, where n is the size of the maximum neighborhood independent set of G and $m = |V|$.
2. V' contains a neighborhood independent vertex set V'' such that $|V''| \geq \frac{n}{2}$.

Corollary 1. ([3]) *There is a constant $\varepsilon > 0$ such that it is **NP**-hard to compute a vertex set $V' \subseteq V$ of a graph $G(V, E)$, with the following properties:*

1. $|V'| \leq nm^\varepsilon$, where n is the size of the maximum neighborhood independent set of G and $m = |V|$.
2. V' contains a neighborhood independent vertex set V'' such that $|V''| = n$.

3 Models

Following Franklin and Wright [6], 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, non-neighbors 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). In our models, 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 perform local computations, and then possibly multicast a message.

Generally there are two kinds of adversaries. A passive adversary (or gossiping adversary) is an adversary who can only observe the traffics through k internal nodes. An active adversary (or Byzantine adversary) is an adversary with polynomial-time bounded computational power who can control k internal nodes. That is, an active adversary will not only listen to the traffics 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 k 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 strongest adversary. Throughout this paper, unless specified otherwise, we will use k to denote the number of nodes that the adversary can control and use n to denote the connectivity of the network.

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 5. (see Franklin and Wright [6])

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 , $\sum_c |\Pr[\text{adv}(M_0, r) = c] - \Pr[\text{adv}(M_1, r) = c]| \leq 2\varepsilon$. 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, $\log \frac{1}{\varepsilon}$ (if $\varepsilon > 0$) and $\log \frac{1}{\delta}$ (if $\delta > 0$).

In order for an adversary to attack the broadcast communication system which is modeled by a strongly n -connected graph, s/he does not need to find all of the n neighborhood disjoint paths $\{p_1, \dots, p_n\}$ in the graph. S/he can choose to control one neighbor vertex on each of the n paths. We therefore give the following definition.

Definition 6. Let G be a strongly n -connected graph, and $P = \{p_1, \dots, p_n\}$ be a witness to the strong n -connectivity of the graph. A set $S = \{v_1, \dots, v_k\}$ of vertices in G is called an eavesdropping vertex set of P if for each path $p_i \in P$ ($i = 1, \dots, n$), there is a $v_j \in S$ which is a neighbor of (at least one of the vertices in) p_i .

Note that in the above definition, k could be larger than or smaller than n .

Remark 1. The problem of finding an eavesdropping vertex set in a strongly k -connected graph is **NP**-hard which will be proved in Section 5.

The following theorem follows straightforwardly from the corresponding theorems in Franklin and Wright [6], and Wang and Desmedt [13].

Theorem 3. (Franklin and Wright [6], Wang and Desmedt [13]) *If A and B are strongly n -connected, and the adversary does not control an eavesdropping vertex between A and B , then there is an efficient $(0, \delta)$ -secure message transmission protocol between A and B .*

4 Reliable and private communication with trapdoors

In this section, we show how to design reliable and private communication systems with trapdoors such that the following condition is satisfied:

- The broadcast communication system modeled by a strongly n -connected graph is robust against a polynomial time bounded k -active adversary where $k \leq cn$ and $c > 1$ is any given constant.

The idea is to use the fact that it is **NP**-hard to find an eavesdropping vertex set of a strongly n -connected graph. It follows that if one designs the graph in such a way that the trusted participants can easily find a witness to the strong n -connectivity of the graph, and the sender and receiver always initiate a communication through this witness, then reliable and private communication is possible. The benefit from using trapdoors in a communication system with broadcast channels is obvious. If we do not use trapdoors then, Franklin and Wright [6]’s results show that a strongly n -connected graph is only robust against k -active adversaries when $k < n$. However, if we use trapdoors in the design of graphs, then with high probability, a strongly n -connected graph is robust against k -active adversaries where $k \leq cn$ and $c > 1$ is any given constant. The reason is that even though the adversary has the power to jam or control $k > n$ vertices in the graph, he does not know which vertices to corrupt such that each path p_i will have a corrupted neighbor (which can eavesdrop on the messages sent through the path p_i).

Definition 7. Let $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ be an ensemble of graphs with the property that each graph in \mathcal{G}_n is strongly n -connected but not strongly $n + 1$ -connected, where \mathcal{N} is the set of positive integers, and let k_n ($n = 1, 2, \dots$) be a sequence of positive integers. The ensemble $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ is called polynomial-time robust against k_n -active adversaries if for every probabilistic polynomial-time algorithm D with the property that for each $G \in \mathcal{G}_n$, $D(G)$ is a k_n -element vertex subset of G , and for every polynomial $p(\cdot)$ and all sufficiently large n , the following condition is satisfied:

If

$$\left(\sum_{\text{graph } G} \text{Prob}(\mathcal{G}_n = G) \right) \neq 0,$$

(that is, with non zero probability there exists a graph G in \mathcal{G}_n , in another words, \mathcal{G}_n is not empty) then the following inequality holds:

$$\left| \sum_{G \in \mathcal{G}_n} \text{Prob}(D(G) \text{ is an eavesdropping vertex set of } G) \right| < \frac{1}{p(n)}.$$

The probabilities in the above definition are taken over the corresponding random variables \mathcal{G}_n and the internal coin tosses of the algorithm D .

Indeed, for the k_n and n in Definition 7, if $k_n < n$, then every ensemble $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ is polynomial-time robust against k_n -active adversaries. So one of the main problems is to design graph ensembles $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ which are polynomial-time robust against k_n -active adversaries for $k_n \geq n$, that is, to design strongly n -connected graphs in which it is hard on the average case to find an eavesdropping vertex set of size $k_n \geq n$. In Section 5, we will outline an approach to generate such kind of graphs. In the remaining part of this section we will demonstrate how to use these graphs to achieve reliability and privacy in broadcast channels.

Protocol I

1. Alice generates a strongly n -connected graph G such that finding a size $k(= cn)$ eavesdropping vertex set is hard, where $c > 1$ is any given constant. (The details will be presented in Section 5).
2. Using a secure channel, Alice sends to the sender and the receiver the n neighborhood disjoint paths $P = \{p_1, \dots, p_n\}$ which is a witness to the strong n -connectivity of G .
3. In order to carry out one communication, the sender and the receiver initiate the communication protocol in Theorem 3 through the n paths in P .

Note that our above protocol is not proactive, that is, it is not secure against a dynamic adversary who after observing one communication will change the vertices he controls. Indeed, it is an interesting open problem to design protocols which are secure against dynamic adversaries.

First assume that Mallory is a k -active adversary where $k = cn$ for some constant $c > 1$, and $P = \{p_1, \dots, p_n\}$ is the set of neighborhood disjoint paths used in Protocol I. Since Mallory does not know how to find a size k eavesdropping vertex set for P (finding such a set is very hard, e.g., as hard as factoring, let say a 1024-bit integer), she does not know which vertices to corrupt so that she can corrupt the system even though she has the power to corrupt $k = cn$ vertices. It follows that the system is robust against a k -active adversary.

5 Strongly n -connected graphs with trapdoors

In this section, we consider the problem of designing strongly n -connected graphs with trapdoors. By using Corollary 1, we will construct practical, average-case hard, strongly n -connected, graphs which are robust against k -active adversaries for $k = n + c$, where c is some given constant. Our following construction is based on the hardness of factoring a large integer and we will not use the approximation hardness results (which will be used to prove theoretical results in the next section).

Construction Let N be a large number which is a product of two primes p and q . We will construct a strongly n -connected graph G with the following property: given the number N and an eavesdropping vertex set for G , one can compute efficiently the two factors p and q . Let x_1, \dots, x_t and y_1, \dots, y_t be variables which take values 0 and 1, where $t = \lfloor \log N \rfloor$. And let $(x_t \dots x_1)_2$ and $(y_t \dots y_1)_2$ denote the binary representations of $\sum x_i 2^{i-1}$ and $\sum y_i 2^{i-1}$ respectively. Then use the relation

$$(x_t \dots x_1)_2 \times (y_t \dots y_1)_2 = N \tag{1}$$

to construct a 3SAT formula C with the following properties (the details of the construction are omitted. Indeed, one can use the constructive proof that 3SAT is **NP**-complete (see, e.g., [8]) to construct the 3SAT formula C though there are more efficient ways for our construction):

1. C has at most $O(t^2)$ clauses.
2. C is satisfiable and, from a satisfying assignment of C , one can compute in linear time an assignment of $x_1, \dots, x_t, y_1, \dots, y_t$ such that the equation (1) is satisfied. That is, from a satisfying assignment of C , one can factor N easily.

Now, by combining the construction in [8] (which constructs a graph GI for each 3SAT formula C with the property that GI has an independent set of size l for some constant $l = O(t^2)$ if and only if C is satisfiable) and the reduction in the proof of Theorem 1, construct a graph $G'(V', E')$ and a number $n = O(t^2)$ with the property that: from a size n neighborhood independent set of G' , one can compute in linear time a satisfying assignment of C . Lastly, the following procedure will generate a strongly n -connected graph G with the property that, from a size $n + c$ eavesdropping vertex set of G , one can compute in linear time a size n neighborhood independent set of G' . Whence from any size $n + c$ eavesdropping vertex set of G , one can compute in polynomial time the primes p and q . As a summary, our construction proceeds as follows.

$$(N, p, q) \rightarrow \text{graph } G' \rightarrow \text{strongly } n\text{-connected graph } G$$

Procedure for generating G from $G'(V', E')$: In the following we construct a multicast graph $f(G') = G(V, E)$ and two nodes $A, B \in V$ (where A denotes the sender and B denotes the receiver) such that there is a neighborhood independent set of size n in G' if and only if A and B are strongly n -connected.

Let $V = \{A, B\} \cup V'$, and $E = E' \cup \{(A, v), (v, B) : v \in V'\}$. It is clear that two paths $P_1 = (A, v_i, B)$ and $P_2 = (A, v_j, B)$ are vertex disjoint and have no common neighbor (except A and B) in G if and only if v_i and v_j have no common neighbor in $G'(V', E')$. Hence there is a neighborhood independent set of size n in G' if and only if A and B are strongly n -connected in G . It is now sufficient to show that from each size $n + c$ eavesdropping vertex set S' of G , one can compute in polynomial time a size n neighborhood independent set of G' .

Since S' is an eavesdropping vertex set of G and G is strongly n -connected, there is at least one size n subset S of S' such that

- S itself is an eavesdropping vertex set of G ;
- S is a neighborhood independent set of G' .

There are $\binom{n+c}{n} = \binom{n+c}{c}$ (which is a polynomial in n) many different size n subsets of S' . Whence by considering all these different size n subset of S' we can compute in polynomial time a size n vertex set S with the above properties.

It is straightforward to see that the above constructed strongly n -connected graph G is robust against k -active adversaries for $k = n + c$ if factoring N is hard, where c is any given constant.

In order to state our main theorem, we need the following assumption of average hardness of factoring.

Hardness Assumption of Factoring: There exists an ensemble $\{X_n\}_{n \in \mathcal{N}}$ (where X_n is a subset of composite numbers of length n) such that for every

probabilistic polynomial-time algorithm D from positive integers to positive integers, every polynomial $p(\cdot)$, and all sufficiently large n , the following condition is satisfied:

$$\left| \sum_{x \in X_n} \text{Prob}(D(x) \text{ is a non-trivial factor of } x) \right| < \frac{1}{p(n)}.$$

Now it is clear that our above discussion implies the following result:

Theorem 4. *Assume the average hardness of factoring, then we can construct a graph ensemble $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ which is polynomial-time robust against k_n -active adversaries, where $k_n = n + c$ for some constant $c > 1$.*

Proof. It follows from the preceding discussions. □

6 Towards theoretical improvements

In the previous section, we outlined a “practical” approach for constructing strongly n -connected graphs which are robust against k -active adversaries for $k = n + c$. In this section we consider theoretical improvements. That is, we will construct strongly n -connected graphs which are robust against k -active adversaries for $k = cn$.

Construction First generate a graph $G'(V', E')$ and a number n which satisfy the conditions of Corollary 1. Secondly, using the method from the previous section of constructing the strongly n -connected graph $G(V, E)$ from $G'(V', E')$, we construct the strongly n -connected graph $G(V, E)$ with the following properties:

1. Two paths P_1 and P_2 in G which go through u_i and u_j respectively are neighborhood disjoint if and only if u_i and u_j have no common neighbor in G' (see the previous section for details).
2. There is a size n neighborhood independent set in G' if and only if there are n neighborhood disjoint paths in G . And from n neighborhood disjoint paths in G one can compute in linear time a size n neighborhood independent set in G' .

From the construction of G from G' , it is straightforward that for any size cn eavesdropping vertex set S' of G , S' contains a size n neighborhood independent set of G' .

By Corollary 1, the graph G is robust against cn -active adversaries.

Remark 2. The above construction shows that, with some reasonable assumption (for example, assume the existence of a probabilistic polynomial time algorithm to generate hard strongly n -connected graphs needed in the above construction, it is possible to construct an ensemble $\{\mathcal{G}_n\}_{n \in \mathcal{N}}$ of strongly n -connected graphs which is robust against polynomial-time k_n -active adversaries, where $k_n = cn$ for some constant $c > 1$.

In this section, we constructed strongly n -connected graphs which are robust against cn -active adversaries. However, these constructions are inefficient and are only of theoretical interests, since the size of the graph G in Corollary 1 will be enormous if we want to make the security of the system to be at least as hard as an exhaustive search of a 1024-bit space. One of the most interesting open questions is how to efficiently generate hard instances of strongly n -connected graphs, especially, for arbitrary number n .

We should also note that, in order to construct the strongly n -connected graphs in this section, we need to construct standard graphs which satisfy the conditions of Corollary 1. That is, we need an algorithm to build graphs whose neighborhood independent sets are hard to approximate in the *average* case. Whence it is interesting (and open) to prove some *average*-case hardness results for the corresponding problems.

Our protocols in this paper are not proactive, that is, not robust against a dynamic adversary who after observing one communication will change the vertices he controls. It is an interesting open problem to design protocols which are secure against dynamic adversaries.

References

1. S. Arora. *Probabilistic Checking of Proofs and Hardness of Approximation Problems*. PhD Thesis, CS Division, UC Berkeley, August, 1994.
2. 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, 1988.
3. M. Burmester, Y. Desmedt, and Y. Wang. Using approximation hardness to achieve dependable computation. In: *Proc. of the Second International Conference on Randomization and Approximation Techniques in Computer Science*, LNCS 1518, pages 172–186, Springer Verlag, 1998.
4. D. Dolev. The Byzantine generals strike again. *J. of Algorithms*, **3**, pp. 14–30, 1982.
5. D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *J. of the ACM*, **40**(1), pp. 17–47, 1993.
6. M. Franklin and N. Wright. Secure communication in minimal connectivity models. In: *Advances in Cryptology, Proc. of Eurocrypt '98*, LNCS 1403, pages 346–360, Springer Verlag, 1998.
7. M. Franklin and M. Yung. Secure hypergraphs: privacy from partial broadcast. In: *Proc. ACM STOC '95*, pages 36–44, ACM Press, 1995.
8. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, San Francisco, 1979.
9. O. Goldreich, S. Goldwasser, and N. Linial. Fault-tolerant computation in the full information model. *SIAM J. Comput.* **27**(2):506–544, 1998.
10. V. Hadzilacos. *Issues of Fault Tolerance in Concurrent Computations*. PhD thesis, Harvard University, Cambridge, MA, 1984.
11. J. Lewis. On the complexity of the maximum subgraph problem. In: *Proc. ACM STOC '78*, pages 265–274, ACM Press, 1978.
12. M. Sudan. *Efficient Checking of Polynomials and Proofs and the Hardness of Approximation Problems*. PhD. thesis, U. C. Berkeley, 1992.

13. Y. Wang and Y. Desmedt. Secure communication in broadcast channels: the answer to Franklin and Wright's question. In: *Advances in Cryptology, Proc. of Eurocrypt '99*, LNCS 1592, pages 443–455, Springer Verlag, 1999.