Privacy from partial broadcast

M. Franklin, M. Yung

Computer Science/Department of Algorithmics and Architecture

Report CS-R9436 June 1994

CWI is the National Research Institute for Mathematics and Computer Science. CWI is part of the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications.
SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a member of ERCIM, the European Research Consortium for Informatics and Mathematics.

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

Privacy from Partial Broadcast

Matthew Franklin*
Moti Yung**

*CWI, P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands
**IBM Research Division, T. J. Watson Center, Yorktown Heights, NY 10598

Abstract

A "partial broadcast channel" allows one processor to send the same message simultaneously to a designated subset of processors. Suppose that a collection of processors are connected by an arbitrary network of partial broadcast channels. We investigate the conditions under which individual processors can exchange messages privately across the network. Our techniques are combinatorial, and do not rely on any complexity theoretic assumptions.

AMS Subject Classification (1991): 94A60
CR Subject Classification (1991): D.4.6
Keywords & Phrases: Cryptography, Privacy, Broadcast.

1. Introduction

This report examines private communication derived from a seemingly "un-private" environment. A collection of processors can communicate among themselves only via "partial broadcasts," i.e., by sending the same message simultaneously from one processor to a subset of processors. The processors can cooperate to execute protocols, but some of them may be dishonest. Given such an environment, can individual processors exchange messages privately?

The answers that we find depend on several factors: (a) the types of partial broadcasts that are allowed; (b) the number and nature of dishonest processors; and (c) the efficiency of communication. We restrict our attention to passive attacks on privacy: No information about message contents is revealed to any (sufficiently small) coalition of processors, assuming all processors execute protocols correctly. The protocols that we present are all extremely simple, have zero probability of error, and guarantee privacy in an information theoretical sense (even when the adversary is all-powerful); our techniques are combinatorial.

An example of a partial broadcast channel is an Ethernet bus or token ring. Any message sent by a processor is heard by all other processors on the bus or ring. An

E-mail addresses: franklin@cwi.nl and moti@research.ibm.com.

Report CS-R9436
ISSN 0169-118X
CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

interconnected collection of busses or rings, with a different subset of processors on each, gives a partial broadcast network.

Another example of a partial broadcast channel is a shared cryptographic key. By publishing an encrypted message, a processor initiates a partial broadcast to the subset of processors that is able to decrypt it. A partial broadcast network arises from a collection of shared keys, e.g., from a key distribution scheme for conferences [3]. A system that uses cryptography to provide for "broadcast encryption" was suggested by Fiat and Naor [7].

Many distributed operation systems support broadcast and group communication primitives. In particular, Dolev and Malki [6] solve basic distributed tasks in an environment supporting broadcast communication. The construction of a broadcast channel from private channels – the Byzantine Agreement Problem [10] – has been a fundamental question in distributed computing theory; our problem can be viewed as its inverse.

The use of private channels as a primitive for general secure computation is shown by Ben-Or, Goldwasser, and Wigderson [2], and Chaum, Crépeau, and Damgård [5], and further explored by many others (see survey [8]). On the other extreme, Goldreich, Goldwasser, and Linial [9] consider the question of secure computation versus active attacks in an environment with only full broadcast (from any processor to all other processors).

We remark that the "radio network" studied by Alon *et al* [1] is similar to one of the partial broadcast networks that we consider (our "neighbor network"). However, a main difference of their model is that a processor receives no messages (i.e., "hears only noise") if it is a recipient of two or more partial broadcasts simultaneously. Their work addresses issues of coordination and scheduling that arise in packet radio networks, and does not consider privacy.

1.1 Motivating Example

Consider a ring of four processors, labeled $1, 2, 3, 4$. Suppose that each processor can partially broadcast to its two neighbors on the ring. The partial broadcast network consists of only these four partial broadcast channels: $1 \rightarrow \{2, 4\}$; $2 \rightarrow \{1, 3\}$; $3 \rightarrow \{2, 4\}$; $4 \rightarrow \{1, 3\}$. Can processor 1 send a message privately to processor 3 over this network?

We first observe that every partial broadcast has either 2 or 4 as a participant. Every message that is sent along every channel is known to at least one of these two processors. Thus it is impossible, via any conceivable protocol, for 1 to send a message to 3 while keeping complete privacy from the coalition $\{2, 4\}$ (unless the sender and receiver share secret information before the start of the protocol, or the coalition of listeners is computationally bounded; our model excludes both of these possibilities). Using the terminology introduced in the next subsection, we say that 2-private communication

from 1 to 3 is impossible from this partial broadcast network.

Is 1-private communication possible from processor 1 to processor 3? At first glance, it appears not, since all message traffic out from 1 can be monitored by a single listener at either 2 or 4. However, no single listener can monitor all traffic coming into 1, and this is crucial.

Consider the following protocol for 1 to send to 3 the bit string m. Processor 2 generates a random bit string r_2 (of the same length as m, from the uniform distribution) and partially broadcasts it to $\{1,3\}$; processor 4 does the same with r_4. Next, processor 1 generates a random bit string r_1 and makes two partial broadcasts to $\{2,4\}$: $r_1 \oplus r_2$, and $m \oplus r_1 \oplus r_4$. These two partial broadcasts are passed on to processor 3 by either (or both) of 2 and 4.

At the end of this protocol, processor 3 has learned the message m, since it has received r_2, r_4, $r_1 \oplus r_2$, and $m \oplus r_1 \oplus r_4$. From the point of view of processor 2, all messages are equally likely (every m is consistent with exactly one choice of r_4); processor 4 also learns no information about m. Thus this protocol is 1-private. In fact, 1-private point to point communication (from *any* sender to *any* receiver) is possible from this particular network. We treat various generalizations of this example in Section 3.3.

1.2 Our Results

In Section 3, we begin with general technical conditions that suffice for an arbitrary hypergraph to support t-private point to point communication (Lemmas 1-3). We then consider "subgraph networks" and "neighbor networks," which arise naturally from an arbitrary undirected graph connecting pairs of processors. We give a complete characterization of privacy from subgraph networks, based on the connectivity of the underlying graph (Theorem 1). For neighbor networks, however, connectivity is not sufficient to determine privacy; we give a partial characterization (Lemma 6), and analyze several interesting examples (Theorems 2-3).

The protocols in Section 3 have communication costs (round and bit complexity) that are polynomial in the length of the message and the size of the network. However, the size of the network might not be polynomial in the number of processors. In Section 4, we restrict our attention to protocols for which communication costs are polynomial in the number of participating processors. We prove lower bounds (Theorems 4-5) and upper bounds (Theorems 6-8) relating t, k, n in the case of a network that allows all partial broadcasts of size k. Our most general upper bound and lower bound, constraining the product tk, are within a log factor of each other: $tk = \tilde{\Theta}(n)$.

2. Model and Definitions

A "partial broadcast" is a message from some sender x to some collection of receivers S. Only the members of S learn the message that was sent by x, while all parties outside of $S \cup \{x\}$ learn nothing about the contents of the message. A "simple channel" from

x to y is a partial broadcast channel from x to the singleton set $S = \{y\}$.

We can represent a partial broadcast channel from x to S as a directed hyperedge connecting x to S. We can represent a collection of partial broadcast channels among n parties $[1 \cdots n]$ as a directed hypergraph on $V = [1 \cdots n]$.

We are interested in using a collection of partial broadcast channels on $[1 \cdots n]$ to achieve a complete network of simple channels on $[1 \cdots n]$. We say that a protocol among the n parties $[1 \cdots n]$ "simulates" a simple channel from x to y, $x, y \in [1 \cdots n]$, if the protocol begins with x choosing a message m, and the protocol ends with y knowing m (with zero probability of error). The protocol may be randomized, i.e., the processors can flip coins during execution.

We say that a simulation is t-private, $1 \leq t \leq n-2$, if no collection $S' \subseteq [1 \cdots n] - \{x, y\}$, $|S'| = t$, learns anything about the message m from the information it receives after honestly participating in the protocol. More formally, the distribution of combined views of any t processors in $[1 \cdots n] - \{x, y\}$ at the end of the protocol should be independent of the message m that was chosen by x. We may consider these t parties to be chosen by a single adversary of unbounded computational power, who is given access to the transcripts of the t parties throughout the protocol. Before the protocol starts, the adversary does not know the particular message that will be sent, or the outcomes of the coin flips of the processors; other than this, the adversary can be assumed to know all relevant information (e.g., specification of protocol, shared information among any processors, inverses of "hard" functions). Notice that this excludes protocols that rely on complexity theoretic assumptions for their privacy. All protocols have zero probability of error, i.e., correct message transmission is guaranteed.

We say that hypergraph H can simulate a complete network of simple channels t-privately if there exists a protocol on H to simulate the simple channel from x to y t-privately for every $x, y \in [1 \cdots n]$. Equivalently, we say that H "supports t-private point to point communication." In this report, we investigate necessary and sufficient conditions on H and t for this task.

It is possible to generalize from the t-private simulation of a simple channel to the t-private simulation of a partial broadcast channel. A partial broadcast channel from x to S is simulated by a protocol if all parties in S know the message at the end of the protocol. The simulation is t-private if no information about the message is revealed to any subset of t processors in $[1 \cdots n] - S - \{x\}$. Given two hypergraphs H, H' on n nodes, and given $t < n$, one could ask whether H can simulate H' t-privately, i.e., whether every partial broadcast channel of H' could be simulated t-privately using only the communication channels of H. We will not explore this generalization in this report.

The following definitions will be useful in the next section. We write a directed hyperedge over a nodeset V as (E^h, E^t) where $E^h \in V$ and $E^t \subseteq V - \{E^h\}$. Definition 1 applies to directed hypergraphs, while Definitions 2 and 3 will only be used for

undirected simple graphs.

Definition 1 *A hypergraph on a nodeset V is "connected" if, for every $x, y \in V$ there exists a sequence of hyperedges $E_1, \cdots, E_k$ and a sequence of nodes $x_0, x_1, \cdots, x_k \in V$ such that (a) $x = x_0$; (b) $y = x_k$; (c) $x_i = E_i^h$ for all i, $0 \leq i \leq k-1$; and (d) $x_i \in E_i^t$ for all i, $1 \leq i \leq k$.*

Definition 2 *The neighborhood of a set S of a graph $G = (V, E)$, denoted $ng(S)$, is the set of all nodes in $V - S$ that are adjacent to some node in S. We may write $\bar{S}$ to denote the set $V - S$.*

Definition 3 *A graph has vertex connectivity k (or is "k-connected") if there are at least k vertex disjoint paths connecting any two nodes. Equivalently, $|ng(S)| \geq k$ for every $S \subset V$, $1 \leq |S| < |V| - k$.*

Definition 4 *A subset $S' \subseteq V$ "cuts" a subset $S \subseteq V$ if S' has nonempty intersection with both S and $V - S$.*

3. Possibility of Private Communication

In this section, we consider the conditions under which t-private point to point communication can be simulated from an underlying network H of partial broadcast channels. We are not concerned in this section with the efficiency of our simulations. In fact, all of our protocols will have round complexity proportional to the diameter of the network, and bit complexity polynomial in the size of the network. However, the size of the network is not necessarily polynomial in the number of processors: A network on n processors can have up to $n2^{n-1}$ distinct partial broadcast channels. In a later section, we treat the case of private simulation using resources that are polynomial in the number of processors.

We begin with some general technical conditions under which t-private communication is possible. Then we consider two specific types of partial broadcast networks: subgraph networks and neighbor networks.

3.1 General Technical Conditions

Three lemmas are proven in this section. Lemma 1 gives a general condition for private communication, and its proof includes a protocol that guarantees privacy whenever privacy is possible. Lemma 2 is a refinement of Lemma 1 that will be useful when subgraph networks are analyzed in Section 3.2. Lemma 3 is a further refinement that will be useful when neighbor networks are analyzed in Section 3.3.

3.1.1 All-Purpose Protocol for Private Communication Our first lemma gives a general condition under which t-private point to point communication is possible from an arbitrary partial broadcast network. The proof relies on a simple protocol for private communication; it is the only protocol used throughout Section 3.

Lemma 1 *Let H be a connected hypergraph on a nodeset V. If, for all $S \subseteq V$, for all $L \subseteq V$ such that $|L| \leq t$, either (a) $S \subseteq L$; or (b) $\bar{S} \subseteq L$; or (c) there exists a hyperedge $E = (E^h, E^t)$ of H such that $(E^t \cup \{E^h\}) \cap L = \emptyset$ and $E^t \cup \{E^h\}$ cuts S, then H supports t-private point to point communication.*

Proof : Suppose that the condition is satisfied. Then the following protocol suffices for any node x to send a message m to any receiver y. Assume that the message m is an element of some finite group.

1. For every hyperedge (E^h, E^t) of H, E^h partially broadcasts $(|E^t|+1)|E^t|$ random messages (group elements chosen according to the uniform distribution), one for each ordered pair $(a, b) \in (E^t \cup \{E^h\})^2$, $a \neq b$. The random message associated with the ordered pair (a, b) is said to have "intended sender" a and "intended receiver" b.

2. Every node finds the sum of messages for which it was the intended receiver and subtracts the sum of messages for which it was the intended sender. If the node is the actual sender x, then it adds to this total the message m that it wishes to send. Call this sum the "final result."

3. Each final result from step 2, except for the final result held by the actual receiver y, is propagated by the nodes openly to the receiver y through a series of partial broadcasts (i.e., without concern for privacy). The sum of all final results, including the final result held by the receiver y, is the message m.

The adversary can learn information about the message only if he can combine final results from some subset $S \subseteq V - \{y\}$, where $x \in S$, with random messages intercepted by some subset $L \subseteq V - \{x, y\}$ such that $|L| \leq t$. Given such an S, L pair, we have that $S \not\subseteq L$ (since $x \in S - L$), and $\bar{S} \not\subseteq L$ (since $y \in \bar{S} - L$). By property (c) of the condition of the Lemma, then, there exists a hyperedge of H such that $(E^t \cup \{E^h\}) \cap L = \emptyset$ and $E^t \cup \{E^h\}$ cuts S. Thus there exists a $z \in S$ and $w \in \bar{S}$ such that $\{z, w\} \subseteq E^t \cup \{E^h\}$. In the protocol, there is a message broadcast for which w is the intended sender and z is the intended receiver. This random message is unrecoverable from the pieces seen by the adversary (since $(E^t \cup \{E^h\}) \cap L = \emptyset$). Since the final result at $z \in S$ is useless to the adversary without knowing this random message, we have a contradiction. □

A somewhat surprising consequence of Lemma 1 is that privacy depends only on the "shape" of the network, and not on its "direction." If H is a directed hypergraph, then

we can define its "directed expansion" $DE(H)$ to include all possible orientations of hyperedges: $DE(H) = H \cup \{(y, S') : (x, S) \in H, y \in S, S' = S + x - y\}$. When H is connected, Lemma 1 implies that t-private point to point communication is possible from $DE(H)$ if and only if it is possible from H. Note that directed expansion arises from the straightforward implementations of partial broadcast by Ethernet or shared conference keys.

3.1.2 Refinements of General Technical Condition The next lemma makes a small modification to the sufficient technical condition from the preceding lemma: L, S can be assumed to be disjoint. This refinement will be useful when subgraph networks are analyzed in Section 3.2.

Lemma 2 *Let H be a connected hypergraph on a nodeset V. If, for all $S \subseteq V$, for all $L \subseteq V - S$ such that $|L| \leq t$, either (a) $\bar{S} \subseteq L$; or (b) there exists a hyperedge (E^h, E^t) of H such that $(E^t \cup \{E^h\}) \cap L = \emptyset$ and $E^t \cup \{E^h\}$ cuts S, then H supports t-private point to point communication.*

Proof : Let S, L contradict the condition of Lemma 1. Then $S - L, L$ contradicts the condition of this claim, as follows. Since S, L contradict part (b) of Lemma 1, $\bar{S} \not\subseteq L$, and thus $\bar{S} \cup L \not\subseteq L$; but $\bar{S} \cup L$ is the complement of $S - L$, so part (a) of Lemma 2 is contradicted by $S - L, L$. Since S, L contradict part (c) of Lemma 1, any hyperedge (E^h, E^t) that is disjoint with L does not cut S. Thus either $E^t \cup \{E^h\} \subseteq S - L$ or $E^t \cup \{E^h\} \subseteq \bar{S} - L$. Since $\bar{S} - L$ is contained in the complement of $S - L$, this means that (E^h, E^t) doesn't cut $S - L$, i.e., part (b) of Lemma 2 is contradicted by $S - L, L$. □

We close with a lemma that will be useful when we consider neighbor networks in Section 3.3. Lemma 3 states that Lemma 2 remains true if S is restricted to be either a single node or a connected union of hyperedges, and of size at most $\frac{1}{2}(n - t)$.

Lemma 3 *Let H be a connected hypergraph on a nodeset V. If, for all $S \subseteq V$ such that S a singleton or a connected union of hyperedges, and such that $|S| \leq \frac{1}{2}(n - t)$, for all $L \subseteq V - S$ such that $|L| \leq t$, either (a) $\bar{S} \subseteq L$; or (b) there exists a hyperedge (E^h, E^t) of H such that $(E^t \cup \{E^h\}) \cap L = \emptyset$ and $E^t \cup \{E^h\}$ cuts S, then H supports t-private point to point communication.*

Proof : Suppose that S, L contradicts the condition of Lemma 2, and $|S| > \frac{n-t}{2}$. Then $V - (S \cup L), L$ is also a counterexample, and $|V - (S \cup L)| \leq \frac{n-t}{2}$.

Suppose that S, L contradicts the condition of Lemma 2, and S is neither a single node of V nor a union of hyperedges of H. Then there exists at least one node $z \in S$ that does not belong to any hyperedge completely contained in S. We claim that

$S - \{z\}, L$ also contradicts the condition of Lemma 2. Otherwise, there is a hyperedge such that $(E^t \cup E^h) \cap L = \emptyset$ and $E^t \cup \{E^h\}$ doesn't cut S and $E^t \cup \{E^h\}$ does cut $S - \{z\}$. This is only possible if z belongs to the hyperedge, and if the hyperedge is completely contained in S, which was assumed impossible for the choice of z. This argument can be repeated until no such z exists, i.e., reducing S to either a single node or a union of hyperedges.

Suppose that S, L contradicts the condition of Lemma 2, and S is not connected. Then S', L is also a counterexample for any connected component S' of S. □

3.2 Subgraph Networks

In this section we derive a necessary and sufficient condition for t-private point to point communication from "subgraph networks." A subgraph network is a directed hypergraph derived from an undirected graph. Let G be an undirected graph on n nodes. The (G, k)-subgraph network is the set of hyperedges such that there is an edge to S from $x \notin S$ if and only if $S \cup \{x\}$ is a connected subgraph of G of size k.

First we show that when H is a (G, k)-subgraph network such that G is $(t+1)$-connected, then Lemma 2 remains true if S is restricted to $|S| = 1$.

Definition 5 *A node $x \in G$ is "t-isolatable" with respect to hypergraph H if there exists a subset of nodes L in G of size t such that no hyperedge (E^h, E^t) that contains x is disjoint with L.*

Lemma 4 *If G is $(t+1)$-connected and no node in G is t-isolatable with respect to the (G, k)-subgraph network, then t-private point to point communication from a (G, k)-subgraph network is possible.*

Proof : Let S, L contradict the condition of Lemma 2, i.e., $L \subseteq V - S$, $|L| \leq t$, $\bar{S} \not\subseteq L$, and no connected subgraph of size k is both disjoint with L and cut by S. Consider the following sequence of sets: $S_0 = S$; $S_{i+1} = S_i \cup \{z_{i+1}\}$ for some $z_{i+1} \in ng(S_i) - L$; and the sequence terminates at S_λ such that either $ng(S_\lambda) \subseteq L$ or $|\bar{S}_\lambda - L| = 1$.

We prove by induction that the condition of the previous lemma is contradicted by each S_i, L, $0 \leq i \leq \lambda$. This will complete the proof, since the condition of the current lemma is contradicted by S_λ, L: either $ng(S_\lambda) \subseteq L$ and thus S_λ contradicts $(t+1)$-connectivity, or $|\bar{S}_\lambda - L| = 1$ and thus $x \in \bar{S}_\lambda - L$ will contradict t-isolatability (see the third paragraph below).

First, we see that $L \subseteq V - S_i$ for all i, $0 \leq i \leq \lambda$. This follows from the fact that $L \subseteq V - S_0$, and that each $z_{i+1} \notin L$.

Second, we see that $\bar{S}_i \not\subseteq L$ for all i, $0 \leq i \leq \lambda$. It is clearly true for $i = 0$. If it is ever false, let j be the smallest index for which $\bar{S}_j \subseteq L$, $1 \leq j \leq \lambda$. Then $|\bar{S}_{j-1} - L| \leq 1$, which by the stopping condition implies that $\lambda \leq j - 1$, a contradiction.

Third, we see that no connected subgraph of size k is both disjoint with L and cut by S_i, for every i, $0 \leq i \leq \lambda$. It is immediately true for $i = 0$. If it is ever false, let j be the smallest index for which there exists a subgraph G' of size k that is disjoint with L and cut by S_j. In particular, G' is not cut by S_{j-1}. This implies that $z_j \in G'$, and also that S_{j-1} is disjoint with G'. Let w be some neighbor of z_j in S_{j-1}. Let v be any node in G' other than z_j whose removal will not disconnect G'. Consider the subgraph $G' - v + w$. It is connected, it has size k, it is disjoint with L, and it is cut by S_{j-1}, a contradiction. □

The next lemma shows that isolatability reduces to connectivity.

Lemma 5 *If G is $(t+1)$-connected, and $t + k \leq n$, then no node in G is t-isolatable with respect to the (G, k)-subgraph network.*

Proof: Suppose, to the contrary, that node x is t-isolatable. Let L be the t locations of G that isolate x, i.e., every connected subgraph of size k that includes x has non-empty intersection with L. Let G' be the largest connected subgraph of G that includes x and excludes all of L. Then $1 \leq |G'| \leq k-1$, else L doesn't fully isolate x. This implies that $V - L - G'$ is non-empty, since $|L| + |G'| \leq t + k - 1 < n$. We must have $ng(G') \subseteq L$, else G' is not maximal. Thus the connectivity of G is at most $|ng(G')| \leq |L| = t$, a contradiction. □

We can now give a complete characterization of private communication from a subgraph network.

Theorem 1 *t-private point to point communication from a (G, k)-subgraph network is possible if and only if $t + k \leq n$ and G is $(t+1)$-connected.*

Proof: If $t+k > n$, then the adversary hears every partial broadcast by occupying any t nodes, and thus private communication is impossible. Suppose that G is not $(t+1)$-connected. Then there exists a non-empty set S such that $V - S - ng(S)$ is non-empty and $|ng(S)| \leq t$. If the nodes of $ng(S)$ are under the control of the adversary, then private communication between any $x \in S$ and any $y \in V - S - ng(S)$ is impossible (else information theoretic two-party secret key exchange is possible).

The reverse direction follows from our last two lemma. □

Using Theorem 1, we can characterize private communication from a "complete k-ary network," i.e., from a partial broadcast network that allows any processor to broadcast to any other $k - 1$ processors. This network will be further studied in Section 4.

Corollary 1 *t-private point to point communication is possible from a complete k-ary network if and only if $t + k \leq n$.*

Proof : If K_n is a complete graph on n nodes, then the (K_n, k)-subgraph network is the complete k-ary network. A complete graph on n nodes is $(n-1)$-connected, and thus Theorem 1 applies for every $t \leq n-2$. □

3.3 Neighbor Networks

In this section, we consider t-private point to point communication from "neighbor networks." Let G be an undirected graph on n nodes. The G-neighbor network is defined to be the set of hyperedges such that there is a hyperedge from x to $ng(x)$ for every vertex x of G. After giving some general technical conditions, we analyze some particular neighbor networks with quite different privacy properties. Although efficiency is not the focus of this section, we point out that the protocol from Lemma 1 is always efficient when executed on a neighbor network, i.e., round complexity and bit complexity are polynomial in the number of processors n and the size of the message (since the size of a neighbor network is linear in the number of processors).

3.3.1 General Technical Conditions for Neighbor Networks The general technical condition of Lemma 3, which holds for all partial broadcast networks, must hold for neighbor networks. However, the analogue of Lemma 4 does not hold for neighbor networks. It is possible that the adversary on a neighbor network can isolate a subset of nodes S, $|S| > 1$, but is unable to isolate any single node. As a concrete example, let H^{2d} be the $2d$-dimensional hypercube, whose nodes are identified with $2d$-bit strings in the standard way. Let G be H^{2d} together with a special node z, where edges connect z to all hypercube nodes of weight $d-1$, d, and $d+1$. Then $S, \{z\}$ is a counterexample to Lemma 3 for the G-neighbor network, where S is all hypercube nodes of weight less than d; but no counterexample satisfies $|S| = |L| = 1$ for Lemma 4. In other words, for this example, a single listener can prevent certain private communications (e.g., from low weight senders to high weight receivers), but no single listener can cut off one node from *all* private communications.

The following lemma refines the sufficient condition of Lemma 3:

Lemma 6 *If the general technical condition of Lemma 3 is violated by the counterexample S, L for a G-neighbor network, then $ng(S) \cup ng(V-S) \subseteq L \cup ng(L)$.*

Proof : For any $x \in ng(S)$, suppose that $x \notin L \cup ng(L)$. Then $x \cup ng(x)$ is disjoint with L, and thus the hyperedge from x to $ng(x)$ is both disjoint with L and cut by S; this is impossible for counterexample S, L. Similarly, the hyperedge from y to $ng(y)$ is both disjoint with L and cut by S if $y \in ng(V-S)$ and $y \notin L \cup ng(L)$. □

As with the subgraph network, it is easy to show that t-private point to point communication from a G-neighbor network requires that G be $(t+1)$-connected. Otherwise, a non-empty subset S can be disconnected from a non-empty subset $V - S - ng(S)$

if the adversary controls the nodes in $ng(S)$, where $|ng(S)| \leq t$. This is the same argument as that of Theorem 1 for subgraph networks.

Unlike the subgraph network, however, this necessary condition is not in general sufficient. In the next few sections, we will illustrate how privacy can vary with respect to connectivity.

3.3.2 Neighbor Networks With Minimal Privacy Although connectivity provides an upper bound on privacy, it does not, in general, imply any lower bound whatsoever. The earlier example of H^{2d} with a special node z demonstrates this: That graph is $(2d)$-connected, while even 1-privacy is unachievable from its neighbor network. An even simpler example is the complete graph K_n, which is $(n-1)$-connected, but cannot support 1-private communication; in fact, *any* node isolates *any* set.

3.3.3 Neighbor Networks With Maximal Privacy In this section, we show that certain graphs G yield neighbor networks with maximal privacy, i.e., t-private communication when G is $(t+1)$-connected.

Example: Let R_n be a ring with n nodes. This graph is 2-connected, and supports 1-private point to point communication between any pair of nodes. Let S be any single node or connected union of hyperedges in the (R_n)-neighbor network, $|S| < \frac{1}{2}n$. Then S is either a single node or m adjacent nodes, $3 \leq m < \frac{1}{2}n$ nodes. Thus $ng(S) \cup ng(V-S)$ contains either three adjacent nodes with central node in S, or two pairs of two adjacent nodes. In either case, it is impossible for $L \cup ng(L)$ to contain all three of these nodes when $|L| = 1$, $L \cap S = \emptyset$. By Lemma 6, then, $|L| \geq 2$.

We now generalize this example to graphs of arbitrary connectivity. We will prove that, for every k, there exists a graph that is k-connected and supports $(k-1)$-private communication. In what follows, we use the term "star" to denote a node and its neighbors; we use the term "upper node" to denote a node in a tree that is neither a leaf nor the parent of a leaf.

Let $T_{k,m}$ be a tree of depth m and degree k, i.e., the root has k children, every internal node has $k-1$ children, and there are m levels excluding the root, $m \geq 2$. Let $G_{k,m}$ be k copies of $T_{k,m}$ with leaves identified ("glued together into super-leaves") so that each super-leaf corresponds to one leaf from each tree, and so that the following property holds: Every pair of super-leaves are siblings in at most one tree. There are many ways to achieve such a mapping.

For example, the mapping of leaves to super-leaves can be created as follows. Let $T^1, \cdots, T^k$ be the k copies of $T_{k,m}$. Choose k primes $p_1, \cdots, p_k$ between k and $k(k-1)^{m-2}$ such that all are relatively prime to $k(k-1)$. Label the leaves of T^i from left to right with increasing multiples of $p_i \bmod k(k-1)^{m-1}$, so that the jth leaf from the left gets the label $jp_i \bmod k(k-1)^{m-1}$. Let all leaves with label i get mapped to the ith super-leaf of $G_{k,m}$. This mapping is a bijection of leaves to super-leaves for

each tree, since each p_i is relatively prime to $k(k-1)^{m-1}$. Suppose that two super-leaves are mapped from siblings in T^i. Then the difference between the labels of the super-leaves must be $\lambda_i p_i \bmod k(k-1)^{m-1}$ for some λ_i, $1 \leq \lambda_i \leq k-2$ (since siblings must be among $k-1$ adjacent leaves in T^i). If the same two super-leaves are mapped from siblings in some other tree T^j, then we would have $\lambda_i p_i = \lambda_j p_j \bmod k(k-1)^{m-1}$, $1 \leq \lambda_i, \lambda_j \leq k-2$. This is impossible, since the two sides of the equality are both less than the modulus, and each contains a prime factor that is missing from the other.

Lemma 7 *$G_{k,m}$ is k-connected.*

Proof : We show that there are k vertex-disjoint paths connecting any pair of nodes x, y. Let T^x denote the tree containing x, and T^y the tree containing y. If u, v are nodes in the same tree, let $lca(u, v)$ denote their least common ancestor. We divide our analysis into cases.

Case 1 is that x, y are in the same copy of $T_{k,m}$ and neither is a super-leaf. Then there is a path from x to y, $k-1$ paths from x to super-leaves $l_1, \cdots, l_{k-1}$, and $k-1$ paths from y to super-leaves $l'_1, \cdots, l'_{k-1}$, such that all are vertex disjoint, and such that all are within that single tree. Then the remaining $k-1$ trees can complete the paths, e.g., a path from l_i to l'_i contained in the ith remaining tree.

Case 2 is that x, y are in the same tree and y is a super-leaf. Then there is a path from x to y, and $k-1$ paths from x to super-leaves $l_1, \cdots, l_{k-1}$, such that all are vertex disjoint, and all within that single tree. Then the ith remaining tree can complete the path from l_i to y.

Case 3 is that x, y are both leaves. Then each of the k trees has its own path from x to y.

Case 4 is that x, y are non-leaves from different trees, and neither is a root. Find two super-leaves l, l' such that exactly one is a descendant of x and exactly one is a descendant of y; this is always possible, no matter what mapping of leaves to super-leaves has been used. There are vertex disjoint paths from x to l and from x to l' entirely within T^x; there are vertex disjoint paths from y to l and from y to l' entirely within T^y. Two vertex disjoint paths from x to y can be constructed from these four partial paths. There are $k-2$ remaining paths from x to leaves and $k-2$ disjoint paths from y to leaves, all disjoint and disjoint with the first two paths. Each of the remaining $k-2$ trees can join pairs of leaves to complete these paths.

Case 5 is that x, y are non-leaves from different trees, and exactly one of them is a root; wlog x is a root. Then there are two vertex disjoint paths in T^x from x to super-leaves l, l', such that l is a descendant of y and l' is not. Then complete the two paths in T^y as in the previous case. Then find the $k-2$ remaining paths as in the previous case.

The remaining case is that x, y are both roots. Then there are two vertex disjoint paths entirely within T^x to super-leaves that have y as their least common ancestor in

T^y. These paths can be completed in T^y. Then find the $k-2$ remaining paths as in the previous two cases. □

Lemma 8 *$(k-1)$-private point to point communication is possible from the $G_{k,m}$-neighbor network.*

Proof : If not, then Lemma 6 implies that there exists S, L such that S is either a single node or a connected union of stars in $G_{k,m}$, $|S| < \frac{1}{2}|V|$, such that $|L| = k-1$, such that L, S are disjoint, and such that $ng(S) \cup ng(V-S) \subseteq L \cup ng(L)$. We derive a contradiction, by a case analysis on the structure of S.

If S is a single node x, then $ng(x)$ contains k elements. For every node $y \in G_{k,m}$, $y \neq x$, we have $|(y \cup ng(y)) \cap ng(x)| \leq 1$. Thus at least k distinct elements are needed in L in order to have $ng(x) \subseteq L \cup ng(L)$. Note that this argument depends on the particular mapping of leaves to super-leaves in $G_{k,m}$. For example, the identity mapping fails because a single super-leaf could be isolated by occupying one parent and one sibling.

Suppose S contains a star centered at a super-leaf, or a star centered at the parent of a super-leaf, but S misses at least one upper node of every tree. Then $ng(S)$ includes at least one upper node from every tree. For any node $y \in L$, $y \cup ng(y)$ can hit at most one of these nodes in $ng(S)$; thus $|L| \geq k$.

Suppose S contains all upper nodes of one of the k trees; call this tree T. Then every super-leaf parent in T is either in S or L; otherwise, some node in $V-S$ is not in $L \cup ng(L)$. For the super-leaf parents in $T \cap S$, every child is either in S or L for the same reason. For the super-leaves in S, every super-leaf parent in every other tree is either in S or L. Adding all of these nodes implies either that S contains more than half the nodes in $G_{k,m}$ or $|L| \geq k$, for sufficiently large m: i.e., $|L| < k$ implies at least $k(k-1)^{m-2} - k$ parents of super-leaves in $T \cap S$, and at least $k(k-1)^{m-1} - k(k-1)$ super-leaves in S, and at least $k(k-1)^{m-1} - k(k-1)$ parents of super-leaves in S from the other $k-1$ trees, for a total of more than $2k(k-1)^{m-1}$ nodes in S; while the size of $G_{k,m}$ is $k(k-1)^{m-1} + k(1 + \sum_{i=0}^{m-2} k(k-1)^i = k(k-1)^{m-1} + k + \frac{k^2}{k-2}((k-1)^{m-1} - 1) < (2 + \frac{2}{k-2})k(k-1)^{m-1}$.

The remaining case is that S is a connected union of stars, but contains no star centered at a super-leaf or centered at the parent of a super-leaf. Then S is fully contained among the upper nodes of a single tree. In this case, there exists an upper node $x \in S$ such that all $k-1$ children of x are in $ng(V-S)$. Then $L \cup ng(L)$ must include all $k-1$ children of x (elements of $ng(V-S)$) and all $(k-1)^2$ children of children of x (elements of $ng(S)$). But for any $y \in L$, $y \cup ng(y)$ can contain no more than $k-1$ of these two generations: It can contain at most one of them if y is in the same tree (since L, S disjoint); and it can contain no more than $k-1$ of them if y is in another tree (maximized when the children of children of x are super-leaves, and y is a

parent of a super-leaf, assuming the worst possible mapping of leaves to super-leaves). Thus $|L| \geq k$ to cover these $k(k-1)$ elements of $ng(S) \cup ng(V-S)$. □

Combining these two lemmas, we have the following theorem:

Theorem 2 *For every $k \geq 2$, there exists a graph G such that G is k-connected and $(k-1)$-private point to point communication is possible from the G-neighbor network.*

3.3.4 Hypercube Neighbor Networks: "Average" Privacy In this section, we analyze neighbor networks in which the underlying graph is a hypercube. Let H^d denote the d-dimensional hypercube, with 2^d nodes and $d2^{d-1}$ edges. Although H^d is d-connected, we prove that t-private communication is possible on the H^d-neighbor network only when $t < \lceil \frac{d+1}{2} \rceil$ (i.e., roughly the average of the minimal and maximal levels of privacy from a d-connected graph). We begin with some facts about neighborhoods of hypercubes. Let $C(x,y)$ denote the number of unordered subsets of size y of a set of size x, i.e., $C(x,y) = \frac{x!}{y!(x-y)!}$.

Fact 1 (Bollobas [4], Theorem 10.5) *Let $S \subset H^d$, where $|S| = \sum_{i=0}^{r} C(d,i)$. Then $|ng(S)| \geq C(d,r+1)$.*

Fact 2 *Let $S, S' \subset H^d$, where $S' = S \cup \{x\}$ for some $x \in H^d$. Then (a) $|ng(S')| \geq |ng(S)| - 1$; and (b) $|ng(S')| + |ng(S' \cup ng(S'))| \geq |ng(S)| + |ng(S \cup ng(S))| - 1$.*

Proof : If $y \in ng(S)$, then $y \notin S$, and there exists $z \in S$ such that y and z are adjacent. Then $y \in ng(S)$ also, unless $y = x$. [Intuitively, everything in the "first shell" around S stays there unless it is "hit" by x (in which case it "falls" into S).] This suffices to prove part (a).

If $y \in ng(S \cup ng(S))$, then $y \notin S \cup ng(S)$, and there exists $z \in ng(S)$ such that y and z are adjacent. Then $y \in ng(S' \cup ng(S'))$ as well, unless $y = x$ or y adjacent to x. Furthermore, $y \in ng(S')$ when y adjacent to x. [Intuitively, everything in the "second shell" around S stays there unless it is "hit" by x (in which case it "falls" into S) or a neighbor is "hit" by x (in which case it "drops" to the "first shell")].

Thus $ng(S) \cup ng(S \cup ng(S)) - \{x\} \subseteq ng(S') \cup ng(S' \cup ng(S'))$. By the triangle inequality, we have $|ng(S')| + |ng(S' \cup ng(S'))| \geq |ng(S)| + |ng(S \cup ng(S))| - |\{x\}| = |ng(S)| + |ng(S \cup ng(S))| - 1$. This proves part (b). □

Fact 3 $|ng(S)| + |ng(S \cup ng(S))| \geq d^2 - d$ *for all* $S \subset H^d$, $|S| = 2$, $d \geq 6$.

Proof : Let $S = \{x, y\}$. By symmetry, the size of the neighborhoods will depend only on the distance from x to y. The relevant neighborhoods of x and y are disjoint when

the distance from x to y is greater than four, so only four cases need to be considered: $\{0^d, 0^{d-1}1\}$, $\{0^d, 0^{d-2}11\}$, $\{0^d, 0^{d-3}111\}$, $\{0^d, 0^{d-4}1111\}$.

When $S = \{0^d, 0^{d-1}1\}$, $ng(S) \cup ng(S \cup ng(S))$ contains $(d-1)$ nodes of weight one (all but one), $\frac{1}{2}d(d-1)$ nodes of weight two (all), and $\frac{1}{2}(d-1)(d-2)$ (all ending in 1). The sum is $d^2 - d$.

When $S = \{0^d, 0^{d-2}11\}$, $ng(S) \cup ng(S \cup ng(S))$ contains d nodes of weight one (all), $\frac{1}{2}d(d-1)-1$ nodes of weight two (all but one), $(d-2)$ nodes of weight three (all ending in 11), and $\frac{1}{2}(d-2)(d-3)$ nodes of weight four (all ending in 11). The sum is $d^2 - d$.

When $S = \{0^d, 0^{d-3}111\}$, $ng(S) \cup ng(S \cup ng(S))$ contains d nodes of weight one (all), $\frac{1}{2}d(d-1)$ nodes of weight two (all), $3(d-3)$ nodes of weight three (all ending in 011,101,110), $(d-3)$ nodes of weight four (all ending in 111), and $\frac{1}{2}(d-3)(d-4)$ nodes of weight five (all ending in 1111). The sum is $d^2 + d - 6$.

When $S = \{0^d, 0^{d-4}1111\}$, $ng(S) \cup ng(S \cup ng(S))$ contains d nodes of weight one (all), $\frac{1}{2}d(d-1)$ nodes of weight two (all), 4 nodes of weight three (all starting 0^{d-4}), $4(d-4)$ nodes of weight four (all ending 0111, 1011, 1101, 1110), $(d-4)$ nodes of weight five (all ending 1111), and $\frac{1}{2}(d-4)(d-5)$ nodes of weight six (all ending 1111). The sum is $d^2 + d - 6$. □

Using these facts, we can prove the following key lemma.

Lemma 9 $|ng(S)| + |ng(S \cup ng(S))| > (d+1)\lceil\frac{d+1}{2}\rceil$, *for all* $S \subset H^d$, $1 < |S| < 2^{d-1}$, $d \geq 7$.

Proof : The case $|S| = 2$ is handled by Fact 3 above. When $2 < |S| \leq d$, let $R \subset S$, $|R| = 2$. By Fact 3, together with Fact 2b, we have $|ng(S)| + |ng(S \cup ng(S))| \geq |ng(R)| + |ng(R \cup ng(R))| - |S - R| \geq d^2 - d + 2 - |S|$. Thus $|ng(S)| + |ng(S \cup ng(S))| \geq d^2 - d + 2 - d > (d+1)\lceil\frac{d+1}{2}\rceil$ for $d \geq 7$.

Otherwise $\sum_{i=0}^{j} C(d,i) \leq |S| < \sum_{i=0}^{j+1} C(d,i)$, for some j, $1 \leq j \leq \frac{d}{2} - 1$. Let $R \subset S$, $|R| = \sum_{i=0}^{j} C(d,i)$. Then $|ng(S)| + |ng(S \cup ng(S))| \geq |ng(R)| + |ng(R \cup ng(R))| - |S - R|$ by Fact 2b. By Fact 1, $|ng(R)| = C(d, j+1) + \Delta$, for some $\Delta \geq 0$. By Fact 1 together with Fact 2a, $|ng(R)| + |ng(R \cup ng(R))| \geq C(d, j+1) + \Delta + C(d, j+2) - \Delta = C(d, j+1) + C(d, j+2)$. Thus $|ng(S)| + |ng(S \cup ng(S))| \geq C(d, j+1) + C(d, j+2) - |S - R| > C(d, j+2) \geq C(d, 3) = \frac{1}{6}d(d-1)(d-2) > (d+1)\lceil\frac{d+1}{2}\rceil$ for $d \geq 7$. □

It is possible to prove a better lower bound by a more careful analysis of hypercube neighborhoods, e.g., using more sophisticated bounds in Bollobas [4]. Nevertheless, the result of the lemma is good enough to prove our desired result:

Theorem 3 *t-private point to point communication from a* H^d*-neighbor network is possible if and only if* $t < \lceil\frac{d+1}{2}\rceil$, *for sufficiently large d.*

Proof : Suppose that $t = \lceil \frac{d+1}{2} \rceil$. Let each node of H^d have a d-bit label in the standard way. The adversary can succeed in isolating the node 0^d by occupying node $0^{d-1}1$ and nodes $0^{2i}110^{d-2i-2}$ for all i, $0 \leq i \leq \lceil \frac{d-1}{2} \rceil - 1$. Then private communication from O^d to any unoccupied node is impossible.

For the other direction, let $S \subseteq V$ be any subset of nodes of H^d, and let $L \subseteq V - S$. By Lemma 3, we can restrict our attention to S a singleton or a union of "stars" of H^d (where a star is a node together with its d neighbors) such that $S < 2^{d-1}$.

The cases where S is a single node or a single star can be handled directly. If $S = \{x\}$, then suppose that L contains k neighbors of x, $0 \leq k \leq d$. Then stars centered at the remaining $d - k$ neighbors of x must be cut by elements of L in $ng(x \cup ng(x))$. But every element of $ng(x \cup ng(x))$ belongs to exactly two stars with centers in $ng(x)$. Thus we need $t \geq k + \lceil \frac{d-k}{2} \rceil \geq 1 + \lceil \frac{d-1}{2} \rceil$. If $S = \{x\} \cup ng(x)$, then suppose that L contains k elements of $ng(S)$, $0 \leq k \leq |ng(S)| = \frac{1}{2}d(d-1)$. Then stars centered at the remaining $\frac{1}{2}d(d-1) - k$ elements of $ng(S)$ must be cut by elements of L in $ng(S \cup ng(S))$. But every element of $ng(S \cup ng(S))$ belongs to exactly three stars with centers in $ng(S)$. Thus we need $t \geq k + \lceil \frac{1}{3}(\frac{1}{2}d(d-1) - k) \rceil \geq 1 + \lceil \frac{1}{3}(\frac{1}{2}d(d-1) - 1) \rceil \geq \lceil \frac{d+1}{2} \rceil$, $d \geq 5$.

When S is a union of more than one star, Lemma 6 implies that the adversary must fail when $|ng(S)| + |ng(V - S)| > |L| + |ng(L)|$. But $|L| + |ng(L)| \leq (d+1)|L|$, so it suffices to show that $|ng(S)| + |ng(V - S)| \geq (d+1)\lceil \frac{d+1}{2} \rceil$. Thus it suffices to show that $|ng(C)| + |ng(C \cup ng(C))| \geq (d+1)\lceil \frac{d+1}{2} \rceil$ for all nodesets C, $1 < |C| < 2^{d-1}$; here C is the set of centers of stars in S. By the preceding lemma, this is true for $d \geq 7$, completing the proof of the theorem. □

The proof of this theorem holds for $d \geq 7$. More detailed analysis yields a proof for smaller d.

4. Efficiency of Private Communication

We say that a simulation is "efficient" if private communication requires only a polynomial overhead in round and bit complexity, with respect to the number of processors n and the length of the message. The protocol from the proof of Lemma 1 in Section 3 is efficient only when the size of the partial broadcast network is polynomial in the number of processors. In this section, we analyze when t-private efficient point to point communication is possible given the ability to make partial broadcasts of size k.

4.1 Lower Bounds

We start with a basic lemma, from which our lower bounds will be derived.

Lemma 10 *t-private point to point communication from partial broadcasts of size k requires that at least $C(n-2,t)/C(n-k,t)$ messages be sent.*

Proof : Suppose that x sends a message to y by participating in a protocol among the n parties wherein a total of l partial broadcasts are sent. Let $\mathcal{S}' = \{S'_1, \cdots, S'_l\}$ denote

the parties involved in these partial broadcasts. Let $\mathcal{S} = \{S_1, \cdots, S_l\}$ denote the parties remaining in the partial broadcasts after x and y are removed from consideration, i.e., $S_i = S_i' - \{x, y\}$. Then $k - 2 \leq |S_i| \leq k$ for all $1 \leq i \leq l$.

Privacy is violated if there exists a set of parties $R \subseteq [1 \cdots n] - \{x, y\}$, $|R| = t$, such that $R \cap S_i \neq \emptyset$ for all $1 \leq i \leq l$. Thus for every R, $|R| = t$, there must exist an $S_j \in \mathcal{S}$ such that

$$(*) \quad R \subseteq \bar{S}_j = [1 \cdots n] - S_j.$$

There are $C(n-2, t)$ choices for R, and security requires that some S_j must satisfy $(*)$ for each possible R. However, each S_i can satisfy $(*)$ for at most $C(n - 2 - |S_i|, t) < C(n - k, t)$ choices of R (i.e., the number of subsets of size t that miss x, y and the parties in $|S_i|$). The lower bound follows. □

Lemma 11 *If t is not a constant, and if $(1 + \alpha)^t$ is polynomial in n and t, then $\alpha = O(\frac{\log n}{t})$.*

Proof : If t is not constant, and if $(1+\alpha)^t$ is polynomial in t, then clearly $\alpha < 1$. Then $2\log(1+\alpha) = \log(1+2\alpha+\alpha^2) > \log(1+\alpha+\frac{\alpha^2}{2-\alpha}) = \log(1+\alpha+\frac{\alpha^2}{2}+\frac{\alpha^3}{4}+\frac{\alpha^4}{8}+\cdots) > \log(1+\alpha+\frac{\alpha^2}{2!}+\frac{\alpha^3}{3!}+\cdots) = \log(e^\alpha) = \alpha$. If $(1+\alpha)^t$ is polynomial in n, t then there exists a $c > 0$ such that $(1+\alpha)^t < n^c$, which implies that $\log(1+\alpha) < \frac{c\log n}{t}$. Combining with the earlier conclusion, we have that $\alpha < \frac{2c\log n}{t}$, i.e., that $\alpha = O(\frac{\log n}{t})$ as required. □

A lower bound for efficient t-private communication follows from these two lemmas.

Theorem 4 *If t is not a constant, then efficient t-private point to point communication from partial broadcasts of size k requires that $kt = O(n \log n)$.*

Proof : Expanding the condition of Lemma 10 (and ignoring constants that cannot affect super-polynomiality), we have $|\mathcal{S}| \geq \frac{n!/(n-t)!t!}{(n-k)!/(n-t-k)!t!} = \frac{n!(n-t-k)!}{(n-t)!(n-k)!} = \frac{n^{\underline{t}}}{(n-k)^{\underline{t}}}$ (using Knuth's notation of $x^{\underline{y}} = x(x-1)\cdots(x-y+1)$). Thus $|\mathcal{S}| \geq (\frac{n}{n-k})^t$ (since $\frac{x}{y} < \frac{x-\lambda}{y-\lambda}$ whenever $x > y$ and $\lambda > 0$). By Lemma 11, if $(\frac{n}{n-k})^t = (1 + \frac{k}{n-k})^t$ is polynomial in n and t, then $\frac{k}{n-k} = O(\frac{\log n}{t})$. That is, there exists a $c > 0$ such that $\frac{k}{n-k} \leq \frac{c\log n}{t}$. Thus $kt \leq c(n-k)\log n < cn \log n$, i.e., $kt = O(n \log n)$. □

Analyzing the result of Lemma 10 in a different way gives a lower bound for the special case when k is very large.

Theorem 5 *If t is not a constant, then efficient t-private point to point communication from partial broadcasts of size k requires that $k \leq n - \Theta(\frac{n}{\log n})$.*

Proof : Towards a contradiction, suppose that $k > n - \frac{c'n}{\log n}$ for all $\epsilon > 0$. This would imply that $(n-k)\log n < \epsilon n < \epsilon kt$ for all $\epsilon > 0$ (the last inequality since t is not constant, and $k > n - \frac{n}{\log n}$). But then $kt > \frac{1}{\epsilon}(n-k)\log n$ for all $\epsilon > 0$, which contradicts the penultimate sentence in the proof of Theorem 4. □

4.2 Upper Bounds for Complete k-ary Networks

Now we give some upper bounds for efficient private communication given a complete k-ary partial broadcast network. In other words, (x, S) is a hyperedge in the network for every $x \in [1 \cdots n]$, for every $S \subseteq [1 \cdots n] - x$, $|S \cup \{x\}| = k$. In fact, all of our protocols for private communication are non-interactive. All messages are initiated by x, and include y as a recipient. We emphasize that Corollary 1 in Section 3.1 considers the same environment as the results of this section, but *without* requiring that communication costs be polynomial in n.

Throughout this section, we assume that the message m is an element of some finite group. We say that $m_1, \cdots, m_l$ are "random additive shares" of m if $\sum_{i=1}^{l} m_i = m$ and $m_1, \cdots, m_l$ are otherwise random. For example, $m_1, \cdots, m_{l-1}$ can be uniformly random elements of the group, and $m_l = m - \sum_{i=1}^{l-1} m_i$.

We start with some simple protocols for when t or k is a constant.

Theorem 6 *t-private point to point communication from a complete k-ary network is efficient whenever k is a constant and $t + k \leq n$.*

Proof : Suppose that x wishes to transmit to y. The message m is split into $l = C(n-2, k-2)$ random additive shares $m_1, \cdots, m_l$. Let $S_1, \cdots, S_l$ be all subsets of $[1 \cdots n] - \{x, y\}$ of size $k-2$. Then x sends m_i to $S_i \cup \{y\}$ for each i, $1 \leq i \leq l$. This scheme is efficient, since the number of partial broadcasts is $C(n-2, k-2) < n^k$, which is polynomial in n when k is a constant. The scheme is private: Each subset of gossipers $L \in [1 \cdots n] - \{x, y\}$ of size t is disjoint with at least one subset $S_i \in [1 \cdots n] - \{x, y\}$ of size $k-2$, and thus misses the additive share m_i. □

Theorem 7 *t-private point to point communication from a complete k-ary network is efficient whenever t is a constant and $t + k \leq n$.*

Proof : Suppose that x wishes to transmit to y. The message m is split into $l = C(n-2, t)$ random additive shares $m_1, \cdots, m_l$. Let $S_1, \cdots, S_l$ be all subsets of $[1 \cdots n] - \{x, y\}$ of size t. Let $\hat{S}_i \subseteq [1 \cdots n] - (S_i \cup \{x, y\})$ be an arbitrary subset of size $k-2$; existence is guaranteed since $t + k \leq n$. Now x sends m_i to $\hat{S}_i \cup \{y\}$, for each i, $1 \leq i \leq l$. This scheme is efficient, since the number of partial broadcasts is $l = C(n-2, t) < n^t$, where t is a constant. The scheme is private, since each subset of gossipers S_i misses the additive share m_i. □

Now we give a more general upper bound result, within a logarithmic factor of the lower bound of Theorem 4.

Theorem 8 *t-private point to point communication from a complete k-ary network is efficient whenever $t + k \leq n$ and $tk = O(n)$.*

Proof : We can assume that t and k are not constant, else one of the preceding theorems can be applied. Thus we can assume that $t < cn$ and $k < cn$ for every $c > 0$.

Suppose that x wants to send a message to y. First x finds a positive integer m such that $tk < (m-1)n$. Then x chooses λm random additive shares of his message, where $\lambda = \lfloor \frac{n-2}{k+m-3} \rfloor$. Now x finds a partition of $[1 \cdots n] - \{x, y\}$ into λ disjoint subsets $S_1, \cdots, S_\lambda$ of size $k + m - 3$ each. Then x partially broadcasts a different share to $S_i^* \cup \{x, y\}$ for all i such that $1 \leq i \leq \lambda$, and for all $S_i^* \subseteq S_i$ such that $|S_i^*| = k - 2$.

We claim that this protocol is efficient. The number of messages sent is $\lambda C(k+m-3, k-2) < \lambda(k+m-3)^{m-1} < n(k+m-3)^{m-1}$ which is polynomial in n and k.

We claim that this scheme is private. The adversary cannot learn any information about the message unless it intercepts every random additive share. Let $S_i' \subseteq S_i$ be the processors in S_i controlled by the adversary. If $|S_i'| < m$, then there exists at least one $S_i^* \subseteq S_i - S_i'$ such that $|S_i^*| = k - 2$, and the adversary cannot intercept the share sent from x to $S_i^* \cup \{x, y\}$. Thus the adversary must control at least m processors in each S_i, $1 \leq i \leq \lambda$. It suffices, then, to show that $m\lambda > t$. Since $tk < (m-1)n$, and since t is not constant, we have that $mn > tk + n > tk + n + (m^2 - m - 3t)$. Since $t, k < cn$ for all $c > 0$, we have that $tk + n + (m^2 - m - 3t) > tk + m(t + k) + (m^2 - m - 3t)$. Thus $mn - mk - m^2 + m > tk + tm - 3t$, and so $t < m\frac{n-k-m+1}{k+m-3} = m(\frac{n-2}{k+m-3} - 1) \leq m\lfloor \frac{n-2}{k+m-3} \rfloor = m\lambda$ as required. □

Can the gap be closed between this upper bound and the lower bound in Theorem 4 for efficient t-private point to point communication?

5. Conclusions and Open Problems

In conclusion, we have analyzed the possibility of achieving private point to point communication from a given network of partial broadcast channels. We give exact characterizations for different sorts of partial broadcast networks, including those derived from subgraphs and neighborhoods of an arbitrary underlying graph. We also consider simulations for which communication costs must be polynomial in the number of processors, and give various upper and lower bounds on relevant parameters for this case.

There are many questions left unexplored, since the space of possible partial broadcast networks is vast. Finding better characterizations for arbitrary partial broadcast networks is a challenging problem, as is the simulation of secure channels from these networks versus stronger adversaries.

One specific direction of future research is to extend our analysis of neighbor networks. What conditions on general G are both necessary and sufficient to allow t-private point to point communication from a G-neighbor network? Notice that G need not be regular, i.e., partial broadcasts can be of different sizes.

There are also natural generalization of neighbor networks. Define a (G, k)-neighbor network to be derived from an undirected graph G as follows: There is a hyperedge to S from x if and only if S are $k-1$ neighbors of x. What conditions on G, k, t allow t-private point to point communication? The same question could be asked for a network that allowed broadcast from any processor x to all nodes in G at distance k or less from x.

ACKNOWLEDGMENTS

The first author would like to thank Annette Bleeker, Ronald Cramer, Ray Hirschfeld, and Berry Schoenmakers for helpful discussions.

REFERENCES

1. N. Alon, A. Bar-Noy, N. Linial, and D. Peleg, "On the complexity of radio communication," ACM STOC 1989, 274–285.
2. M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," ACM STOC 1988, 1–9.
3. C. Blundo, A. De Santis, A. Herzberg, S. Kutten, U. Vaccaro, and M. Yung, "Perfectly-secure key distribution for dynamic conferences," Crypto 1992, 471–486.
4. B. Bollobas, *Combinatorics*, Cambridge University Press, 1986.
5. D. Chaum, C. Crépeau, and I. Damgård, "Multiparty unconditionally secure protocols," ACM STOC 1988, 11–19.
6. D. Dolev and D. Malki, "On distributed algorithms in a broadcast domain," ICALP 1993, 371–387.
7. A. Fiat and M. Naor, "Broadcast encryption," Crypto 1993.
8. M. Franklin and M. Yung, "Varieties of secure distributed computing", *Proc. Sequences II, Methods in Communications, Security and Computer Science*, Positano, Italy, June 1991, 392–417.
9. O. Goldreich, S. Goldwasser, and N. Linial, "Fault-tolerant computation in the full information model," IEEE FOCS 1991, 447–457.
10. L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," ACM Trans. on Programming Lang. and Systems (1982), 382–401.