Global Broadcast by Broadcasts Among Subsets of Players

Matthias Fitzi[1] Ueli Maurer[1]

Abstract — **In the standard model with only pairwise communication channels, unconditionally secure broadcast among n players is achievable if and only if the number t of corrupted players satisfies $t < \frac{n}{3}$. We show that, when additionally given broadcast among each subset of three players then global broadcast is achievable if $t < \frac{n}{2}$.**

I. Introduction

Given a set $P = \{p_1, \ldots, p_n\}$ of n players, the goal of *broadcast* is to let some distinct player $d \in P$ (called *dealer*) reliably distribute a value to all players in P, i.e., all correct (i.e. uncorrupted) players must receive the same value (*agreement*), and if the dealer is correct then this must be the value the dealer intended to distribute (*validity*).

In this paper we focus on broadcast protocols that are unconditionally secure against an adversary that may actively corrupt up to t of the n players. To actively corrupt a player means to make him deviate from the protocol in an arbitrarily malicious way. Unconditionally secure means that the correctness of the protocol does not rely on any further restriction on the power of the adversary than the threshold t of players he can corrupt during the protocol.

Since the network typically consists only of communication channels among subsets of players and some of the players, especially the dealer, may be corrupted by the adversary, broadcast is a non-trivial problem.

Pease, Shostak, and Lamport [2] proved that, according to the standard communication model of a complete synchronous network of pairwise authentic channels among each pair of players, unconditionally secure broadcast is achievable if and only if $t < \frac{n}{3}$. The communication model considered in this paper extends this standard model by a synchronous network of authentic broadcast for each subset $S \subseteq P$ of the players of cardinality $|S| = 3$, i.e.,

- for every subset of three players and for any selection of a dealer among them there is a broadcast channel, and
- for every such channel, all involved players are authentic, i.e., every correct player is able to assign a received message to its corresponding broadcast invocation.

A broadcast primitive or protocol for n players that is secure against t corrupted players is called *(n,t)-broadcast*.

II. Results

Theorem 1 *Given $(3,1)$-broadcast, $(n, \lfloor \frac{n-1}{2} \rfloor)$-broadcast is achievable for any $n \geq 3$.*

The basic idea is to take some known broadcast protocol (e.g [2]) for some *virtual* player set Q ($|Q| = n'$) in the standard model that tolerates $t' < \frac{n'}{3}$ corrupted players among Q — where, for the moment, n' can be supposed to arbitrary, i.e., not necessarily dependent on k. Instead of letting the virtual players directly participate in the protocol, every virtual player q_i is simulated by some specific collection $S_i \subseteq P$ of the *actual* players (according to player simulation in [1]). If it can be achieved that at most $t' < \frac{n'}{3}$ players q_i are incorrectly simulated then the protocol achieves broadcast among the players in Q (with respect to the players q_j that are correctly simulated). Finally, broadcast among the players in P can then easily be derived from broadcast among the players in Q.

The following proposition immediately follows from [1].

Proposition 1 *A player $q_i \in Q$ of any protocol among a player set Q can be simulated correctly by a collection of players S if broadcast among the players in S is possible and less than $\frac{|S|}{2}$ players in S are corrupted.*

Proof of Theorem 1: The proof of this theorem is based on a recursive construction that, for any $k > 0$, allows to achieve $(2k+3, k+1)$-broadcast from $(2k+1, k)$-broadcast. Finally, $(3,1)$-broadcast can then be used as a base for the recursive construction in order to achieve any $(n, \lfloor \frac{n-1}{2} \rfloor)$-broadcast.

Let P be a set of $2k+3$ players and assume $(2k+1, k)$-broadcast to be achievable among any $S \subset P$ with $|S| = 2k+1$. We define a set Q of $n' = \binom{2k+3}{2k+1}$ virtual players and involve them in some standard broadcast protocol that tolerates $t' < \frac{n'}{3}$ player corruptions. We now let every possible collection $S \subset P$ of $|S| = 2k+1$ players from P simulate exactly one player $q_i \in Q$. Such a player q_i is simulated correctly if at least $k+1$ of the simulating players are correct themselves (since $k+1$ constitutes a majority and hence broadcast among S works correctly and hence Proposition 1 applies), i.e., at least $k+1$ of the simulating players in S must be corrupted by the adversary in order to corrupt the corresponding virtual player. Hence at most $t' \leq \binom{k+1}{k+1}\binom{k+2}{k}$ players of the original protocol can be corrupted which are given by all simulating collections $S \subset P$ of cardinality $|S| = 2k+1$ including at least $k+1$ corrupted players. Since $k > 0$ we get

$$\frac{n'}{t'} = \frac{\binom{2k+3}{2k+1}}{\binom{k+2}{k}} = \frac{2(2k+3)}{k+2} = 4 - \frac{2}{k+2} > 3,$$

and hence strictly less than a third of the players in the original protocol is corrupted. Finally we can let every simulated player send his result to every simulating player who then can compute the outcome of the broadcast by a majority voting on all received values. ∎

References

[1] M. Hirt and U. Maurer, "Complete characterization of adversaries tolerable in secure multi-party computation," *Proc. 16th ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 25–34, Aug. 1997.

[2] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *Journal of the ACM*, 27(2):228–234, Apr. 1980.

[1]Department of Computer Science, ETH Zürich, CH-8092 Zurich, Switzerland. E-mail: {fitzi,maurer}@inf.ethz.ch. Research supported by the Swiss National Science Foundation, SPP project no. 5003-045293.