Brief Announcement: Towards Byzantine Broadcast in Generalized Communication and Adversarial Models

Chen-Da Liu-Zhang
Department of Computer Science, ETH Zurich, Switzerland
lichen@inf.ethz.ch

Varun Maram
Department of Computer Science, ETH Zurich, Switzerland
vmaram@inf.ethz.ch

Ueli Maurer
Department of Computer Science, ETH Zurich, Switzerland
maurer@inf.ethz.ch

Abstract

Byzantine broadcast is a primitive which allows a specific party to distribute a message consistently among n parties, even if up to t parties exhibit malicious behaviour. In the classical model with a complete network of bilateral authenticated channels, the seminal result of Pease et al. [6] shows that broadcast is achievable if and only if $t < n/3$. There are two generalizations suggested for the broadcast problem – w.r.t. the adversarial model and the communication model. Fitzi and Maurer [2] consider a (non-threshold) *general adversary* that is characterized by the subsets of parties that could be corrupted, and show that broadcast can be realized from bilateral channels if and only if the union of no three possible corrupted sets equals the entire set of n parties. On the other hand, Considine et al. [1] extend the standard model of bilateral channels with the existence of b-minicast channels that allow to locally broadcast among any subset of b parties; the authors show that in this enhanced model of communication, secure broadcast tolerating up to t corrupted parties is possible if and only if $t < \frac{b-1}{b+1}n$. These generalizations are unified in the work by Raykov [5], where a tight condition on the possible corrupted sets such that broadcast is achievable from a complete set of b-minicasts is shown.

This paper investigates the achievability of broadcast in *general networks*, i.e., networks where only some subsets of minicast channels may be available, thereby addressing open problems posed in [4, 5]. Our contributions include: 1) proposing a hierarchy over all possible general adversaries for a clean analysis of the broadcast problem in general networks, 2) showing the infeasibility of a prominent technique – used to achieve broadcast in general 3-minicast networks [7] – with regard to higher b-minicast networks, and 3) providing some necessary conditions on general networks for broadcast to be possible while tolerating general adversaries.

1 Motivation

To the best of our knowledge, current works on the achievability of broadcast in general networks [7, 4] focus on the problem of Byzantine agreement for the concrete case of 3-minicast channels, and mainly against a threshold adversary in the range $n/3 \le t < n/2$. We continue the line of research w.r.t. general b-minicast channels. We remark that – as noted

in [1] – when $b > 3$, perfectly secure broadcast can be realized even when there is no honest majority, in contrast to Byzantine agreement.

2   Models

Notation. Let $P = \{P_1, \ldots, P_n\}$ be a set of n parties. We say that a list $\mathcal{S} = (\mathcal{S}_0, \ldots, \mathcal{S}_{k-1})$ is a k-partition of P if $\bigcup_{i=0}^{k-1} \mathcal{S}_i = P$ and all $\mathcal{S}_i$ and $\mathcal{S}_j$ are pair-wise disjoint. In addition, we denote the set of parties from P *minus* the two sets $\mathcal{S}_i$ and $\mathcal{S}_j$ from the k-partition $\mathcal{S}$ as $\mathcal{S}_{\downarrow i,j} \coloneqq P \setminus (\mathcal{S}_{i \bmod k} \cup \mathcal{S}_{j \bmod k})$.

Adversary. We assume the existence of a central adversary that corrupts a subset of parties at the onset of a protocol execution. Corrupted parties are *Byzantine*, i.e. can behave in an arbitrary way. We consider a general adversary structure $\mathcal{A}$ [3], which specifies the possible subsets of parties that the adversary can corrupt. We require that $\mathcal{A}$ be monotone, i.e., $\forall a, a'$ $(a \in \mathcal{A})$ and $(a' \subseteq a) \implies a' \in \mathcal{A}$. In this paper, we are interested in adversary structures which satisfy the k-chain condition [5].

▶ **Definition 1.** *An adversary structure $\mathcal{A}$ is said to contain a k-chain if there exists a proper k-partition $\mathcal{S} = (S_0, \ldots, S_{k-1})$ of the party set P such that $\forall i \in [0, k-1]$ $\mathcal{S}_{\downarrow i,i+1} \in \mathcal{A}$. An adversary structure is k-chain-free if it does not have a k-chain.*

Communication Network. A general network $\mathcal{N}$ among a set of parties $\mathcal{P}$ is a monotone[1] set of subsets of P. Given a general network $\mathcal{N}$, we have $\{P_{i_1}, \ldots, P_{i_k}\} \in \mathcal{N}$ if and only if there is a partial broadcast channel that allows broadcast to be realized locally among $\{P_{i_1}, \ldots, P_{i_k}\}$ – such a channel is also known as *k-minicast channel* [5]. As instantiations, the classical model with bilateral channels [6] corresponds to the network structure $\mathcal{N}$, where $\mathcal{N}$ contains all possible subsets of P with size 2; the complete b-minicast model [5] is a network structure which contains all partial broadcasts of size at most b.

3   Our Results

We extend results for general 3-minicast networks to general b-minicast networks and address open questions posed in both of the papers [4, 5], i.e. to study broadcast achievability in general communication models where only a subset of b-minicast channels may be available.

Hierarchy of Adversary Structures. We propose a hierarchy of adversary structures based on the chain terminology introduced in [5]. This allows us to analyze the feasibility of broadcast in smoothly evolving minicast models in a meaningful way. Recall that in the complete b-minicast communication model, broadcast tolerating adversary structure $\mathcal{A}$ is achievable if and only if $\mathcal{A}$ is $(b+1)$-chain-free [5]. Let the *weakest* adversary class be $\mathfrak{A}^{(0)} = \{\mathcal{A} \subseteq 2^P \mid \mathcal{A}$ is 3-chain-free$\}$ where broadcast is achievable with only bilateral channels, and the *strongest* adversary class be $\mathfrak{A}^{(n)} = \{\mathcal{A} \subseteq 2^P \mid \mathcal{A}$ contains an n-chain$\}$ where broadcast is not possible among the n parties unless we assume a global broadcast primitive in the first place. The subsequent classes of adversary structures in-between are defined as: $\forall b \in [3, n-1], \mathfrak{A}^{(b)} = \{\mathcal{A} \subseteq 2^P \mid \mathcal{A}$ contains a b-chain and is $(b+1)$-chain-free$\}$. One can order the adversary classes as follows: $\mathfrak{A}^{(0)} \leq \mathfrak{A}^{(3)} \leq \mathfrak{A}^{(4)} \leq \ldots \leq \mathfrak{A}^{(n)}$. This forms a partition over all adversary structures, since for $b > 3$, any $\mathcal{A} \in \mathfrak{A}^{(b)}$ also contains an implicit $(b-1)$-chain (and lower). Observe that given an adversary $\mathcal{A} \in \mathfrak{A}^{(b)}$, broadcast tolerating $\mathcal{A}$ is impossible in the complete $(b-1)$-minicast model, but is achievable in the

---

[1] If $N \in \mathcal{N}$ and $N' \subseteq N$ then $N' \in \mathcal{N}$.

complete b-minicast model. This allows us to study the b-minicast channels that play an essential role in realizing broadcast against this particular adversary. We do not consider stronger classes, e.g. $\mathfrak{A}^{(b+1)}$, because [5] shows that broadcast is not achievable under these adversaries even if all b-minicast channels are available. Also, regarding weaker classes such as $\mathfrak{A}^{(b-1)}$, we know that there already exist protocols that implement secure broadcast in the complete $(b-1)$-minicast model without any b-minicast channel.

Naive Emulation of Virtual Parties. We discuss the limitations of an application of the *virtual emulation* technique to construct broadcast protocols in a (possibly incomplete) b-minicast network. The technique involves generating a new set of *virtual parties* V which are emulated by the original party set P. For example, [7] addresses the problem of achieving broadcast in a general 3-minicast network by using the available 3-minicasts to emulate virtual parties, thereby reducing the original problem to that of implementing broadcast among $P \cup V$ in the underlying communication model with bilateral channels that is secure against an *extended* adversary structure (to account for corrupted virtual parties). We show that this kind of strategy is not applicable for general b-minicast channels since we deal with significantly stronger adversaries. More concretely, if $\mathcal{A}_P$ is an adversary structure with respect to P that contains a b-chain and is $(b+1)$-chain-free (i.e., $\mathcal{A}_P \in \mathfrak{A}^{(b)}$), we prove that – irrespective of the subset of b-minicasts available for emulation in the communication model, the extended adversary $\mathcal{A}_{P \cup V}$ contains a b-chain, and thus, there does not exist any protocol among $P \cup V$ that achieves secure broadcast in the reduced $(b-1)$-minicast model.

▶ **Lemma 2.** *For $b > 3$: given an adversary structure $\mathcal{A}_P \in \mathfrak{A}^{(b)}$, for any possible set of virtual parties V emulated using b-minicast channels in the communication model, the corresponding extended adversary $\mathcal{A}_{P \cup V}$ contains a b-chain.*

Essential Partial Broadcasts. We identify some *types* of b-minicast channels that are essential for secure broadcast to be possible against general adversaries. The following characterization of partial broadcast channels also allows us to derive a lower bound on the number of b-minicasts required for the parties to broadcast globally in any general network.

▶ **Theorem 3.** *Secure broadcast on a general network $\mathcal{N}$ tolerating any general adversary $\mathcal{A} \in \mathfrak{A}^{(b)}$ is possible only if: for every b-chain in $\mathcal{A}$, namely $\mathcal{P} = (\mathcal{P}_0, \ldots, \mathcal{P}_{b-1})$, there is a b-minicast channel in $\mathcal{N}$ that has non-empty intersection with the sets $\mathcal{P}_0, \ldots, \mathcal{P}_{b-1}$.*

───── **References** ─────

1   J. Considine, M. Fitzi, M. Franklin, L. A. Levin, U. Maurer, and D. Metcalf. Byzantine agreement given partial broadcast. *Journal of Cryptology*, 18(3):191–217, July 2005.

2   M. Fitzi and U. M. Maurer. Efficient byzantine agreement secure against general adversaries. In S. Kutten, editor, *DISC*, volume 1499 of *LNCS*, pages 134–148. Springer, 1998.

3   Martin Hirt and Ueli Maurer. Complete characterization of adversaries tolerable in secure multi-party computation. In *PODC*, volume 97, pages 25–34, 1997.

4   A. Jaffe, T. Moscibroda, and S. Sen. On the price of equivocation in byzantine agreement. In D. Kowalski and A. Panconesi, editors, *ACM PODC '12*, pages 309–318. ACM, 2012.

5   Raykov P. Broadcast from minicast secure against general adversaries. In M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann, editors, *ICALP 2015*, volume 9135 of *LNCS*, pages 701–712. Springer, Berlin, Germany, 2015.

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

7   D. V. S. Ravikant, M. Venkitasubramaniam, V. Srikanth, K. Srinathan, and C. P. Rangan. On byzantine agreement over (2,3)-uniform hypergraphs. In R. Guerraoui, editor, *DISC 2004*, volume 3274 of *LNCS*, pages 450–464. Springer, 2004.