

On Necessary Conditions for Secure Distributed Computation

Rafail Ostrovsky*

Moti Yung†

Abstract

What assumptions are required to achieve an *unconditionally secure* distributed circuit evaluation in a fully connected network? This question was addressed with respect to the allowed number of malicious players [BGW, CCD, RB], given that *every channel* is unconditionally secure. In this paper we investigate whether the security of all channels is also a necessary condition.

[BGW, CCD] showed how secure computation can be achieved, provided that a constant fraction of the total number of players is honest. An insecure channel can be modeled as faults on both ends of the channel. Thus, as long as the number of such “faulty” players is smaller than the fraction established in [BGW, CCD], the channels can be made insecure. However, an insecure channel seems to be a much weaker fault than a corruption of both players. Thus, can a bigger fraction of insecure channels be tolerated? In this paper we show that this is not the case. That is, we show that in some cases the perfect security of multi-party protocols in a fully connected network requires *all* the channels to be physically secure. In particular, we show a simple protocol (for three parties) for which if privacy of even one channel is compromised, the protocol can not be computed securely. Thus, we establish that the security of *all* channels is not only sufficient (by the work of [BGW, CCD]), but also *necessary*. The lower bound holds even if players follow the protocol. That is, we establish our impossibility result even if all the players are honest but curious — if they follow the protocol exactly, but try to extract additional information “on the side”. Thus, our result gives a pure security perspective of the impossibility. An additional feature of our result is its extreme simplicity, which is usually hard to come by for the lower bound proofs.

AMS(MOS) Subject Classification: 68M10, 68P25, 68Q05.

*MIT Lab. for Computer Science Cambridge, MA 02139. E-mail to: “raf@theory.lcs.mit.edu”. Part of this work was done while the author was at the IBM Research, T.J. Watson Research Center, Yorktown Heights, NY 10598.

†IBM Research, T.J. Watson Research Center, Yorktown Heights, NY 10598. E-mail to: “moti@ibm.com”.

1 Introduction

This paper deals with feasibility results concerning the implementation of unconditionally secure computation in insecure communication environments. That is, we examine the feasibility of multi (≥ 3)-party secure computation. We concentrate on global computations in which all parties compute a private (possibly random) output.

The question of secure distributed computation received a lot of attention over the past decade, which culminated in the work of [GMW1] where they showed a way to compute any poly-size function on a fully-connected network of processors securely, under some general cryptographic assumptions, provided that more than $\frac{2}{3}$ of the processors are honest.

The cryptographic assumptions were then eliminated in the work of [BGW, CCD] where it was established that if every two processors can communicate secretly, one can achieve secure computation without any cryptographic assumptions for three or more processors (provided that either more than $\frac{1}{2}$ are honest while the rest are honest but may be curious, or that more than $\frac{2}{3}$ of the processors are honest in the case the rest of the processors may be Byzantine (malicious)). However, both the work of [BGW, CCD] and further extensions by [RB, BG] require each pair of processors to have a secure communication channel. In this work we examine whether this condition can be weakened, and provide a strong negative answer to this question.

That is, we show a gap between a network with physically secure channels and a network without such security measures. To do so, we exhibit a protocol for which there is no perfectly secure implementation on the second model even when only one channel is unprotected. The first model, on the other hand, is known to be universal for perfectly secure computations (even when up to $\frac{1}{3}$ of participants are malicious [BGW, CCD].) This shows formally that adding physical security to *all* channels is not just sufficient but also *necessary*.

Notice that any insecure channel can be made secure using suitable cryptographic assumptions [GMW2, GHY]. The resulting protocol, however, is only as secure as the cryptographic assumption which was utilized. Instead, we are interested in the question of absolute security, independent of any assumptions. Our proof establishes that:

MAIN THEOREM: *Providing physical security to all channels is NECESSARY to the achievement of perfect security in distributed multi-party secure computation.*

Thus, our result justifies the model of physically secure channels as a model which achieves universality in the set of perfectly secure computation even when the parties are computationally unlimited, provided that *every* channel is secure. It also justifies the use of cryptography (and achieving only computational security) when such channels are not available.

The rest of the paper is organized as follows: In section 2 we describe our model. Section 3 explains our proof, while in section 4, we review recent and related results.

2 The model

We consider two models of computation for multi-party protocols. Both have computationally unlimited users in a fully-connected network. The first one has secure channels between each pair of users, while in the second model all (or some) channels are unprotected.

Both models have been used in various contexts in the past. For example, Feldman and Micali [FM] implemented a fast Byzantine Agreement protocol in the secure channel model and left an open problem whether such a fast Byzantine Agreement protocol (even with relaxed performance) can be implemented on the insecure channel model. (They also show a simulation of the private channel model in the insecure channel model using cryptographic assumptions, but here we are interested in perfect security.)

A bit more formally, we consider the model of multi-party protocols which is the standard system of communicating machines [GMR]. Each player is a probabilistic Turing machine with a private computation environment. They share communication tapes and communicate by writing messages on these tapes in a synchronous fashion.

Each pair of machines share a communication tape (for each such tape only two machines are allowed to write on it). When each machine has “read access” to all communication tapes this is the *insecure channels* model (or the board model), on the other hand when each of the communication tapes can be accessed only by the pair of parties which are allowed to write on it – this is the *secure channels* model. A natural intermediate models with some private and some insecure channels can be defined as well.

A point worth mentioning is that the notion of secure channels can be somewhat relaxed. That is, the channels do not have to be totally secure, as we can adapt the wire-tap model of Wyner [W]. (He basically assumes certain rate of being caught by the eavesdropper. By using linear codes according to the rate one gets that the users get the message while the eavesdropper does not.) Thus, the notion of “secure” channel can be interpreted in the above, weaker sense.

3 A simple proof of our result

We start with a review of Shamir, Rivest and Adleman’s impossibility of Mental Poker [SRA]. This is a basic impossibility result (which seems to have been somewhat forgotten!) It shows that two players cannot deal a secret, disjoint, and random card hands based on information theory and open communication.

Shamir, Rivest and Adleman proof considers the minimal non-trivial scenario of two un-faulty (but possibly “curious”) players, A and B , who try to deal a hand of one card to each out of a deck of three cards $\{x, y, z\}$. (This scenario can be extended to larger decks and any size of hands). Let the protocol *execution* M be the finite sequence of transmitted messages. M should coordinate the cards drawn by each player (the cards drawn are a function of M and the player’s internal computations). The hands should be drawn uniformly at random and be disjoint. Let x (y) be the actual card received by A (B).

Let S_A (S_B) be the set of candidate cards A (B) could have gotten, given M . S_A cannot contain only x since then B can compute the hand of A by simulating all possible computations of A consistent with the execution. S_A cannot contain all three cards since B will not be able to get any card disjoint to S_A : regardless of what card B gets, there is a computation consistent with M in which A may get the same card. Thus, S_A consists of two cards. Similarly, S_B must contain two cards. The total size of both candidate sets is four while the total size of the deck is three, thus there is an intersection between S_A and S_B , and there must be a

computation consistent with M in which both players gets the same card (in this case it is z). This contradicts the disjointness requirement which implies the impossibility.

We can now give an overview of our proof: consider the case of three nodes in the network. Note that if three players want to play and one channel is suspected not to be secure, they could use cryptography to solve the problem [GMW2, GHY]. But this solution relies on unproven complexity assumptions, while we are interested in retaining *perfect* (i.e. information-theoretic) security. A protocol for dealing cards (Mental-Poker) for three or more players which achieved information theoretic security was implemented, provided that all channels are secure in 1983 [BaFu]. The basic idea of our proof is to show that a three player Mental-Poker is impossible if the channels are not secure.

We consider dealing of cards where there are three players and four cards in the deck. We start by assuming totally open communication. Given an execution (the protocol message sequence), none of the players can have only one card in the set of cards which are candidate to be taken to his hand consistent with the execution — this will violate security. The case of more than three cards in the candidate set is impossible as well (since if one player has more than three possible cards, then another player must have a non-disjoint card in his candidate set violating disjointedness of hands). Thus all players have two cards in their candidate set. However, again the deck is too small to provide disjointedness of the candidate sets, that is, there must be a computation consistent with the execution in which hands are not disjoint. Thus, the dealing protocol is impossible when all channels are insecure.

Can we make the result sharper and consider a network with only one insecure channel (while the rest are secure)? We assume the same scenario and problem as before with an intermediate model of a single insecure channel (say between A and B , while both channels to C are private). We have the execution which is the three sequences of messages over the three channels M_{ab} , M_{ac} , M_{bc} .

Again, in this case as before, no player may have a candidate set bigger than three. C may have a candidate set of size one. However, both A and B must have candidate sets of size two, since otherwise C who knows all the messages can determine their cards. The sum of the sizes of the candidate card sets from C 's point of view is at least 5 while only 4 cards are present in the deck. (The other players see even less information than C and thus their view of the computation should also leave candidate set of size 2.) This is a contradiction to disjointedness.

Thus, in this case even one public channel prevents a perfectly secure implementation:

Theorem: *There exists protocols which can not be executed securely in the information-theoretic sense on a fully-connected network if any one of the channels is compromised.*

This proves the necessity of secure channels in the case of computationally unbounded parties.

4 Related work

Recently, [OVY] considered a two-party asymmetric games when one of the players is infinitely-powerful while the other is polynomially-bounded. Using the proof method similar to the above, they were able to show that information-theoretic Oblivious Transfer protocol is impossible to achieve. Moreover, non-interactive Oblivious Transfer was also shown to be impossible. On the positive side, they were able to show that if one-way functions exist, then any two-party asymmetric game (for example, Oblivious Transfer protocol) is possible to implement.

Further study of requirements for multi-party secure computation, when the network is not fully connected was done by [DDWY], where they presented tight results on the required connectivity of the network in order to preserve security.

Acknowledgments

The authors thank Silvio Micali and Joan Feigenbaum for helpful discussions on the topic of the paper.

References

- [BaFu] I. Barany, and Z. Furedy, *Mental Poker with Three or More Players*, Info. and Cont. v. 59, 1983, pp. 84-93.
- [BGW] Ben-Or M., S. Goldwasser and A. Wigderson, *Completeness Theorem for Noncryptographic Fault-tolerant Distributed Computing*, STOC 1988, ACM, pp. 1-10.
- [BG] Beaver D., S. Goldwasser *Multiparty Computation with Faulty Majority* FOCS 1989, IEEE, pp. 468-473.
- [CCD] D. Chaum, C. Crepeau and I. Damgard, *Multiparty Unconditionally Secure Protocols*, STOC 1988, ACM, pp. 11-19.
- [DDWY] D. Dolev, C. Dwork, O. Waarts and M. Yung, *Secure Message Transmission*, FOCS 1990, IEEE.
- [FM] P. Feldman and S. Micali, *Optimal Algorithms for Byzantine Agreement*, STOC 1988, ACM, pp. 148-161.
- [GHY] Z. Galil, S. Haber and M. Yung, *Cryptographic Computations and the Public-Key Model*, The 7-th Crypto 1987, Springer-Verlag, pp. 135-155.
- [GMW1] S. Goldreich, S. Micali and A. Wigderson, *Proofs that Yields Nothing But their Validity*, FOCS 1986, IEEE, pp. 174-187.
- [GMW2] S. Goldreich, S. Micali and A. Wigderson, *How to Play any Mental Poker*, STOC 1987, ACM, pp. 218-229.
- [GMR] S. Goldwasser, S. Micali and C. Rackoff, *The Knowledge Complexity of Interactive Proof-Systems*, STOC 1985, ACM, pp. 291-304.

- [OVY] R. Ostrovsky, R. Venkatesan, M. Yung *On The Complexity of Asymmetric Games*, manuscript.
- [RB] T. Rabin and M. Ben-Or, *Verifiable Secret Sharing and Multiparty Protocols with Honest Majority*, STOC 1989, ACM, pp. 73-85.
- [SRA] A. Shamir, R. Rivest and L. Adleman, *Mental Poker*, Technical Memo MIT (1979).
- [W] Wyner, A.D., *The Wire Tap Channel* Bell System J., 54, 1981, pp. 1355-1387.