On the Secrecy-Complexity of Distributed Computation of a Binary Function of Non-uniformly Distributed Variables

Eytan Modiano and Anthony Ephremides
Electrical Engineering Department
University of Maryland
College Park, MD

This paper examines the communication complexity of secure distributed computation. This is a relatively new area in the field of secret communications, which studies the message exchange process that is needed to support a task such as a distributed computation or the implementation of a protocol, under secrecy constraints. This problem in its simplest form was introduced in [2], and it involved two processors, P_x and P_y which were interested in exchanging the values of two variables X and Y. P_x knows the value of X, and P_y knows the value of Y and they communicate according to a predetermined protocol. The objective is for both processors to compute the value of a boolean function F(X,Y). Originally the problem focused on the determination of the minimum number of bits that need to be exchanged to pursue this computation. In [1] a secrecy aspect was introduced by considering an eavesdropper who monitors the channel in order to obtain information about the value of F(X,Y), and who is aware of the protocol used by the processors. The communicating processors wish to keep the probabilities of F(X,Y)=1, before and after the communication takes place, e-close to each other for the eavesdropper; i.e.

$$|Pr\{F(X,Y)=1\,|I\} - Pr\{F(X,Y)=1\}| \leq \epsilon$$

where I denotes the observed values of the quantities communicated between P_x and P_y. To achieve this objective some of the bits are transmitted over a private secure channel that is inaccessible to the eavesdropper. The problem is to find a protocol that minimizes the worst case number of bits needed to be transmitted securely in order for the above inequality to hold. It was shown in [1] that when X and Y are uniformly distributed over their range, for any boolean function F and $\epsilon > 0$, P_x and P_y need to exchange no more than $2log(1/\epsilon) + 16$ bits securely. The effects of noise on the above protocol were studied in [3] where it was shown that in the presence of noise in the eavesdropper's channel the transmission of fewer secure bits is required.

In this paper we study the effects of the non-uniform distribution of X and Y on the number of bits which must be transmitted over the secure channel in order for the security requirement to be satisfied. The performance of the protocol developed in [1] depended heavily on the fact that P_x and P_y were uniformly distributed over their range. We need therefore to propose an alternative protocol which is not as sensitive to the distribution of X and Y. Our protocol is based on the partitioning of the function table of F(X,Y), which is, however, significantly different from the one used in [1]. As before, the two processors exchange partitioning information one bit at a time. Initially P_x sends one bit to P_y indicating to which of a possible two partitions of the function table X belongs, then P_y responds by sending one bit to P_x, further partitioning the function table. This process continues until the values of X and Y (and consequently F(X,Y)) are obtained by both processors. The partitions of the function table based on the distribution of F(X,Y) are formed as follows:

Let, P(x,y) be the Probability Density Function of X and Y which takes values in $\mathcal{X} \times \mathcal{Y}$. The first partition of $\mathcal{X}$ into $\mathcal{X}'$ and $\mathcal{X}''$ is accomplished by first splitting $\mathcal{X}$ into m_1 subsets, $\mathcal{X}^1, \mathcal{X}^2, \ldots, \mathcal{X}^{m_1}$ where,

$$\mathcal{X}^i = \left\{x : \frac{m_1 - i}{m_1} \leq P(F=1\,|\,x) < \frac{m_1 - i + 1}{m_1}\right\}$$

and then dividing the elements of each of the $\mathcal{X}^i$s between $\mathcal{X}'$ and $\mathcal{X}''$ So as to approximate as closely as possible the condition $P(F=1|\mathcal{X}') = P(F=1|\mathcal{X}'')$. Upon receiving the partitioning information from P_x, P_y proceed in forming it's first partition of $\mathcal{Y}$ using a similar procedure to the one outlined above. We show that, when P_x and P_y proceed in this way, the size of the function table, a_i, after i partitions is bounded by:

$$a_i \leq \frac{n}{2^i} + \sum_{j=0}^{i-1} \frac{m_j}{2^{i-j}}$$

and, the probability, P_i, of the function F being equal to 1 after i partitions is bounded by:

$$P_i \leq P_0 + 3K \sum_{j=0}^{i-1} \left\lceil \frac{2^{j+1}}{n - \sum_{l=0}^{j-1} m_l 2^l} \right\rceil + \sum_{l=0}^{i-1} \frac{2}{m_l}$$

where, P_0 is the prior probability of the function F being equal to 1, m_i is the number of subsets used in forming the i^{th} partition, n is the cardinality of the ranges of X and Y(assumed to be the same), and K is defined by:

$$K = \frac{\max_{x,y} P(x,y)}{min_{x,y} P(x,y)}.$$

In order to determine the worst case number of secret bits that our protocol requires we must minimize the table size, a_i, subject to the constraint that $P_i \leq P_0 + \epsilon$. The resulting minimum, a_{min}, is the smallest table size possible that does not violate the security requirement. The actual number of secret bits transmitted in the worst case is given by $\lceil log(a_{min}) \rceil$. Preliminary results indicate that this number is independent of n, and is proportional to $log(1/\epsilon)$ and $log(K)$.

References

[1] A. Orlitsky and A. El Gamal, "Communication With Secrecy Constraints," *Proc. of the 16th Annual ACM Symposium on the Theory of Computing*, Atlanta, GA, April 1984, pp. 217-224.

[2] A.C. Yao, "Some Complexity Questions Related to Distributed Computing," *Proc. of the 11th Annual Symposium on the Theory of Computing*, Washington D.C., May 1979, pp. 209-213.

[3] E. Modiano and A. Ephremides, "Communication Complexity of Secure Distributed Computation in the Presence of Noise," Submitted to *IEEE Transactions on Information Theory* (also presented at the *1990 IEEE International Symposium on Information Theory*, San Diego, CA, January 1990).