Set Reconciliation with Nearly Optimal Communication Complexity

Yaron Minsky
Dept. of Computer Science
Cornell University
yminsky@cs.cornell.edu

Ari Trachtenberg
Electrical and Computer Engineering
Boston University
trachten@bu.edu

Richard Zippel
The Interdisciplinary Center
Herzliya, Israel *and*
Compaq Cambridge Research Lab
rz@alum.mit.edu

Abstract — **We consider the problem of efficiently reconciling two similar sets held by different hosts while minimizing the communication complexity. This type of problem arises naturally from gossip protocols used for the distribution of information, but has other applications as well. We describe an approach to such reconciliation based on the encoding of sets as polynomials. The resulting protocols exhibit tractable computational complexity and nearly optimal communication complexity. Moreover, these protocols can be adapted to work over a broadcast channel, allowing many clients to reconcile with one host based on a single broadcast.**

I. Introduction

Gossip protocols [1] spread information through a network of hosts by random contacts between pairs of hosts. Through many such uncoordinated exchanges, information is spread throughout the system. The information disseminated by a gossip protocol usually consists of a set of distinct entries. When a pair of hosts exchange information, they must reconcile their respective data sets. What makes this reconciliation difficult is that the hosts do not know *a priori* which data elements need to be transmitted, and which are already known by the other host.

We formalize the problem of reconciling two hosts' data sets as follows: given a pair of hosts A and B, each with a set of length-b bit-strings, S_A and S_B respectively, how can each host determine the mutual difference between the two sets with a minimal amount of communication? We call this the *set reconciliation* problem.

II. The Basic Protocol

We associate with each set $S = \{x_1, x_2, \ldots, x_n\}$ a *characteristic polynomial* $\chi_S(Z)$, defined to be:

$$\chi_S(Z) = (Z - x_1)(Z - x_2)(Z - x_3)\cdots(Z - x_n)$$

Note that the zeros of $\chi_S(Z)$ are exactly the elements of S. Thus, S can be computed from $\chi_S(Z)$ be factoring.

Let $\Delta_A = S_A \setminus S_B$, *i.e.*, the set of elements in S_A but not in S_B, and let $\Delta_B = S_B \setminus S_A$. Let m be $|\Delta_A| + |\Delta_B|$, the size of the symmetric difference between S_A and S_B.

1. Hosts A and B evaluate $\chi_{S_A}(Z)$ and $\chi_{S_B}(Z)$ respectively at the same $\overline{m}$ sample points, where $\overline{m}$ is an upper bound on m.
2. The evaluated values from A and B are combined to compute $\chi_{S_A}(Z)/\chi_{S_B}(Z)$ at the sample points. Note that
$$\frac{\chi_{S_A}(Z)}{\chi_{S_B}(Z)} = \frac{\chi_{\Delta_A}(Z)}{\chi_{\Delta_B}(Z)},$$
since terms corresponding to elements in $S_A \cap S_B$ cancel out. Because the sum of the degrees of the numerator and denominator of of $\chi_{\Delta_A}(Z)/\chi_{\Delta_B}(Z)$ is m, the $\overline{m}$ sampled values can be interpolated to recover the coefficients of this reduced rational function.
3. By factoring $\chi_{\Delta_A}(Z)$ and $\chi_{\Delta_B}(Z)$, the elements of Δ_A and Δ_B are recovered.

So as to minimize communication and computation costs, all calculations are performed over a finite field $\mathbb{F}_q$, for $q \geq 2^b$.

III. Extensions

The algorithm outlined above assumes that the reconciling hosts have a (close) bound $\overline{m}$ on the size of the symmetric difference between their two sets. In the absence of a close bound on m, the hosts need to detect that enough samples have been taken to properly interpolate the rational function $\chi_{\Delta_A}(Z)/\chi_{\Delta_B}(Z)$. One approach is to attempt several reconstruction of the rational function from different random evaluated samples. A probabilistic analysis gives a relation between the the number of identical reconstructions found and the probability of error ϵ.

Our approach to set reconciliation can be adapted for use over a broadcast channel, as follows. Host A broadcasts a message with $\overline{m}$ samples of $\chi_{S_A}(Z)$. A receiving host with a set sufficiently close to S_A can then compute Δ_A. If A includes extra values from random sample points, a receiving host can determine with high probability whether or not it has obtained computed Δ_A correctly. Proofs as well as more complete references can be found in [2].

Acknowledgments

We are grateful to Ramin Takloo-Bighash, Edward Reingold, Fred Schneider, Lazar Trachtenberg, and Alexander Vardy for stimulating discussions. The implementation was developed with the help of Eyal Adler.

References

[1] A.J. Demers, D.H. Greene, C. Hause, W. Irish, and J. Larson, "Epidemic algorithms for replicated data maintenance," in *Proc. 6th Annual ACM Symp. on Princ. of Distr. Comp.*, Vancouver, Canada, pp. 1-12, 1987.

[2] Y. Minsky, A. Trachtenberg, and R. Zippel, "Set reconciliation with nearly optimal communication complexity," Tech. Rep. TR2000-1813, Cornell University, 2000.