CCIT Report #821 November 2012

DISTRIBUTED SPARSE SIGNAL RECOVERY FOR SENSOR NETWORKS

Stacy Patterson, Yonina C. Eldar, and Idit Keidar

Department of Electrical Engineering
Technion - Israel Institute of Technology, Haifa, Israel
{stacyp,yonina,idish}@ee.technion.ac.il

ABSTRACT

We propose a distributed algorithm for sparse signal recovery in sensor networks based on Iterative Hard Thresholding (IHT). Every agent has a set of measurements of a signal x, and the objective is for the agents to recover x from their collective measurements at a minimal communication cost and with low computational complexity. A naïve distributed implementation of IHT would require global communication of every agent's full state in each iteration. We find that we can dramatically reduce this communication cost by leveraging solutions to the distributed top-K problem in the database literature. Evaluations show that our algorithm requires up to three orders of magnitude less total bandwidth than the best-known distributed basis pursuit method.

***Index Terms*—** compressed sensing, distributed algorithm, iterative hard thresholding, top-K

1. INTRODUCTION

In compressed sensing, a sparse signal $x \in \mathbb{R}^N$ is sampled and compressed into a set of M measurements, where M is typically much smaller than N. If these measurements are taken appropriately, then it is possible to recover x from this small set of measurements [1].

Compressed sensing is an appealing approach for sensor networks, where measurement capabilities may be limited due to both coverage and energy constraints. Recent works have demonstrated that compressed sensing is applicable to a variety of sensor networks problems including event detection [2], urban environment monitoring [3] and traffic estimation [4]. In these applications, measurements of the signal are taken by sensors that are distributed throughout a region. The measurements are then collected at a single fusion center where signal recovery is performed. Due to limits in bandwidth, storage, and computation capabilities, it may be more efficient, and sometimes even necessary, to perform signal recovery in the network in a distributed fashion.

Distributed solutions for compressed sensing have begun to receive attention lately. For example, one work proposes a distributed subspace pursuit recovery algorithm for a mixed-support set model [5]. This work assumes that every agent knows the sensing matrix of every other agent. The need for global knowledge of these matrices presents a scalability bottleneck as individual sensors do not have the capacity to store and process a large number of these matrices. Several works have proposed distributed basis pursuit algorithms for sparse signal recovery in sensor networks where the measurement matrices are not globally known [6, 7, 8]. In these algorithms, agents collaborate to solve a convex relaxation of the original recovery problem. Each agent stores its own estimate of the signal x, and, in each iteration, it updates this estimate based on communication with its neighbors in the network. This approach requires that every agent solve a local convex optimization problem in each iteration. While these algorithms use only local communication, each agent must send its entire estimate vector to every neighbor in every iteration. This vector is not necessarily sparse until the algorithm converges, and therefore, the messages can be quite large. As a result, these algorithms have a large total bandwidth cost. Furthermore, simulations show that this bandwidth cost increases dramatically as the network connectivity increases.

We propose an alternative approach to distributed sparse signal recovery in sensor networks that is based on *Iterative Hard Thresholding* (IHT) [9]. In our distributed implementation of IHT, which we call D-IHT, all agents store identical copies of the estimate of x. In each iteration, every agent first performs a *local computation* to derive an intermediate vector. The agents then perform a *global computation* on their intermediate vectors to derive the next iterate. A naïve distributed implementation of IHT would require global communication of the intermediate vector of each agent in every iteration. We find that we can dramatically reduced the communication cost of this global computation by leveraging solutions to the distributed top-K problem in the database literature [10, 11, 12]. Our evaluations show that D-IHT requires up to three orders of magnitude less total bandwidth than the best-known distributed basis pursuit method. D-IHT is also computationally simpler since it does not require that agents solve local convex optimization problems. While, in this work, we present

The work of Y. Eldar is supported in part by the Israel Science Foundation under Grant no. 170/10, and in part by the Ollendorf Foundation. This work was funded in part by the Arlene & Arnold Goldstein Center at the Technion Autonomous Systems Program, a Technion fellowship, and an Andrew and Erna Finci Viterbi Fellowship.