Title:
Distributed Optimization Algorithms for Networked Systems

Presenter:
Michael M. Zavlanos

Abstract:
Distributed optimization methods allow to decompose an optimization problem into smaller, more manageable subproblems that are solved in parallel. For this reason, they are widely used to solve large-scale problems arising in areas as diverse as wireless communications, optimal control, machine learning, artificial intelligence, computational biology, finance, and statistics, or problems with a separable structure that are amenable to distributed implementations. In this talk we present the Accelerated Distributed Augmented Lagrangians (ADAL) algorithm, a novel decomposition method for optimization problems that involve a separable convex objective function subject to convex local constraints and linear coupling constraints. Optimization using Augmented Lagrangians (ALs) combines the low computational complexity of first order optimization methods with fast convergence speeds due to the regularization terms included in the AL. In its centralized version, optimization using ALs is an excellent general purpose method for constrained optimization problems and enjoys a large amount of literature. However, decentralized methods that employ ALs are few, as decomposition of ALs is a particularly challenging task. We establish convergence of ADAL and show that it has a worst-case $O(1/k)$ convergence rate. Moreover, we show that ADAL converges to a local minimum of the problem when the objective function is non-convex, and that it can handle uncertainty and noise in which case it generates sequences of primal and dual variables that converge to their respective optimal sets almost surely. We provide numerical simulations for wireless network optimization problems that suggest that the proposed method outperforms the state-of-the-art distributed Augmented Lagrangian methods that are known in the literature. Moreover, we present a Random Approximate Projections (RAP) method for decentralized optimization problems with SDP constraints. Unlike other methods in the literature that employ Euclidean projections onto the feasible set, our method is computationally inexpensive as it relies only on subgradient steps in the direction that minimizes the local constraint violation. We show that the algorithm converges almost surely and can also handle inexact problem data. We demonstrate our approach on a distributed estimation problem involving networks of mobile sensors estimating a set of hidden states that are governed by linear dynamics up to a user-specified accuracy.