Large-Scale Parallel Simulations of Distributed Detection Algorithms for Collaborative Autonomous Sensor Networks

Presented at CASIS

Anton Yen, Peter Barnes, Bhavya Kailkhura,Priyadip Ray,
Deepak Rajan, Katie Schmidt, Ryan Goldhahn

May 23, 2018

LLNL-PRES-749650

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Why Distributed Detection?

- Cost – lots of cheap sensors rather than one big one
- Reliability – wider coverage area, can suffer attrition
- Performance – approaches that of centralized with proper choice of a test statistic

Take Measurements at Each Sensor

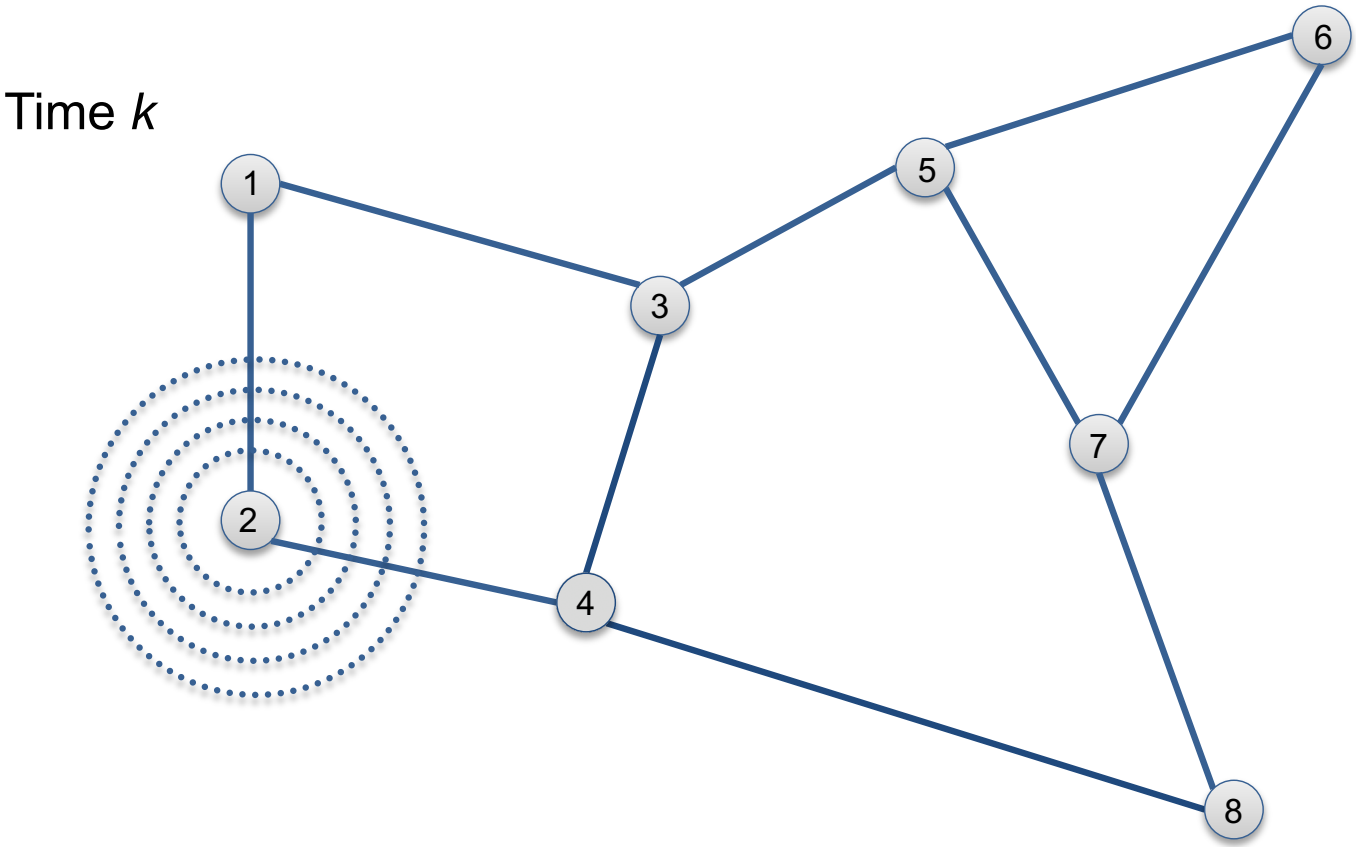

Form Local Test Statistics

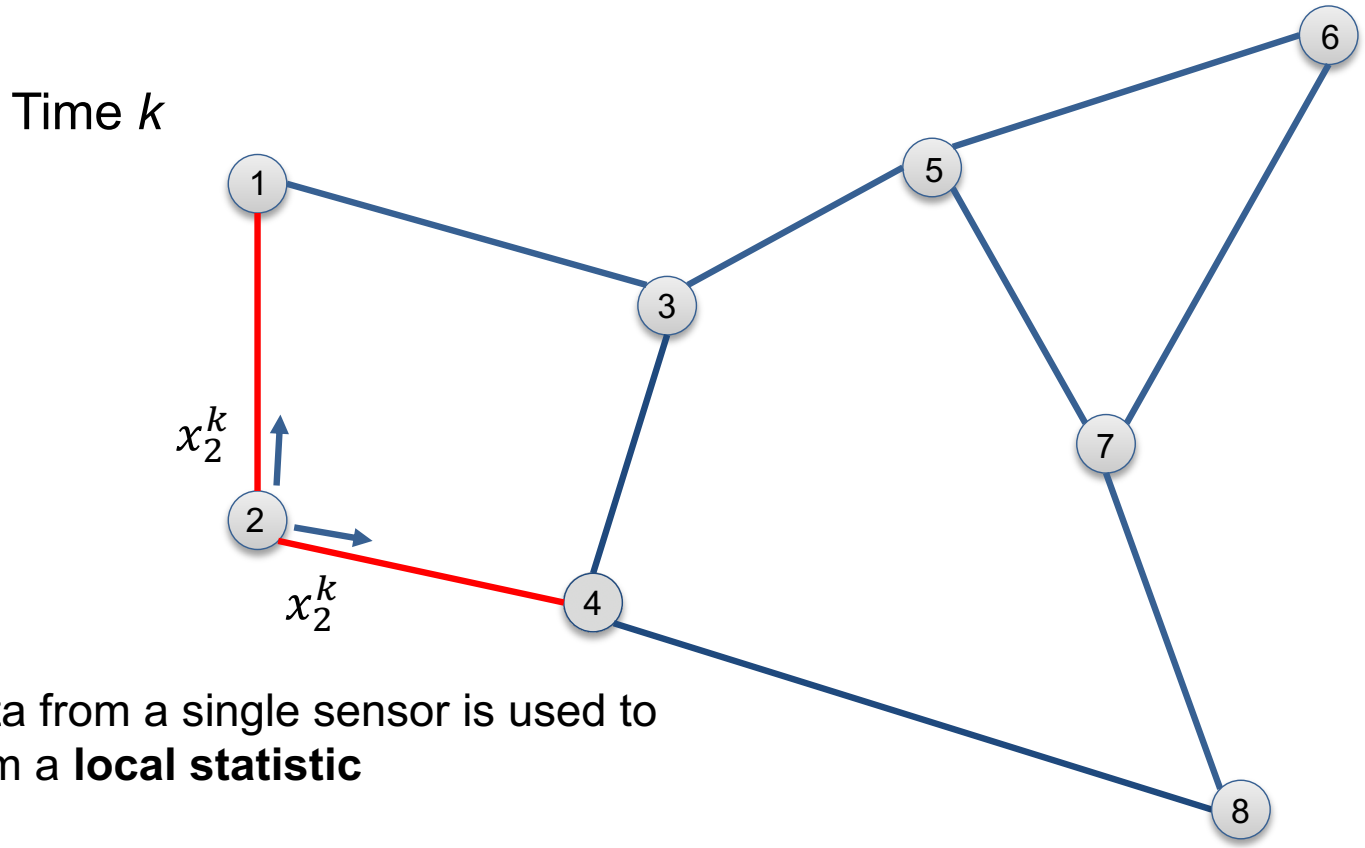

Data from a single sensor is used to form a **local statistic**

Fuse Neighbor State Information

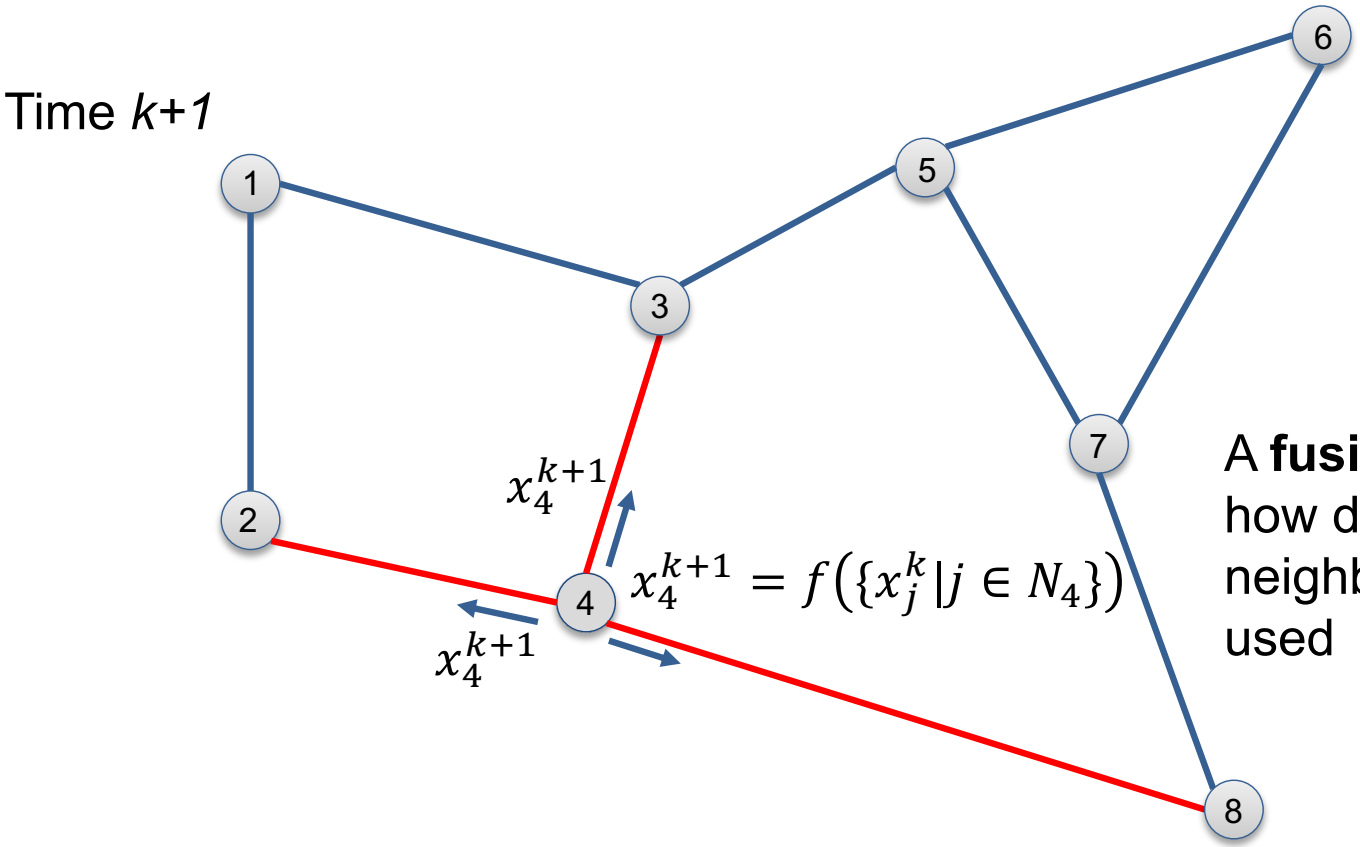

A **fusion rule** defines how data from neighboring nodes is used

Propagate Information Across the Network

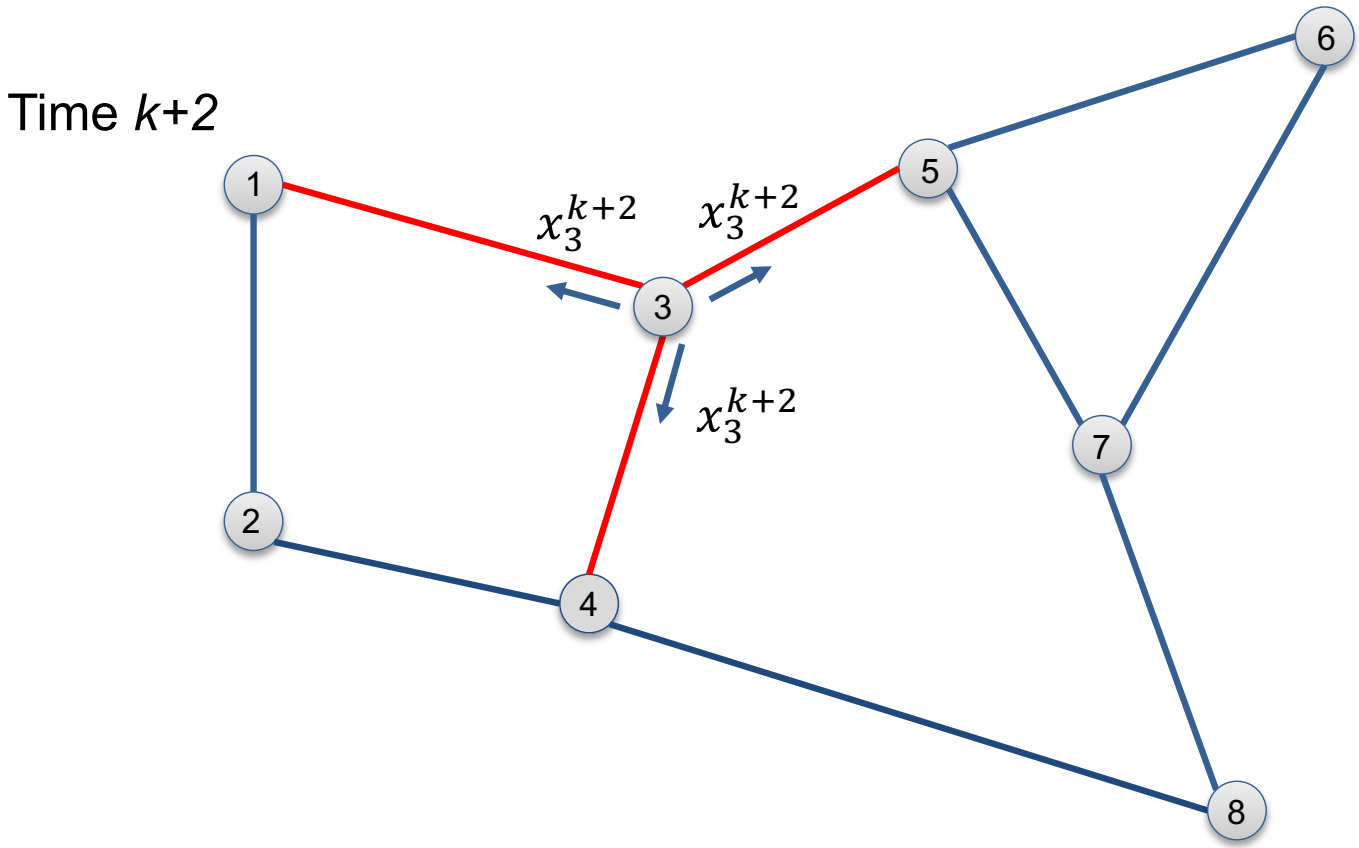

Trade Latency for Communications Overhead

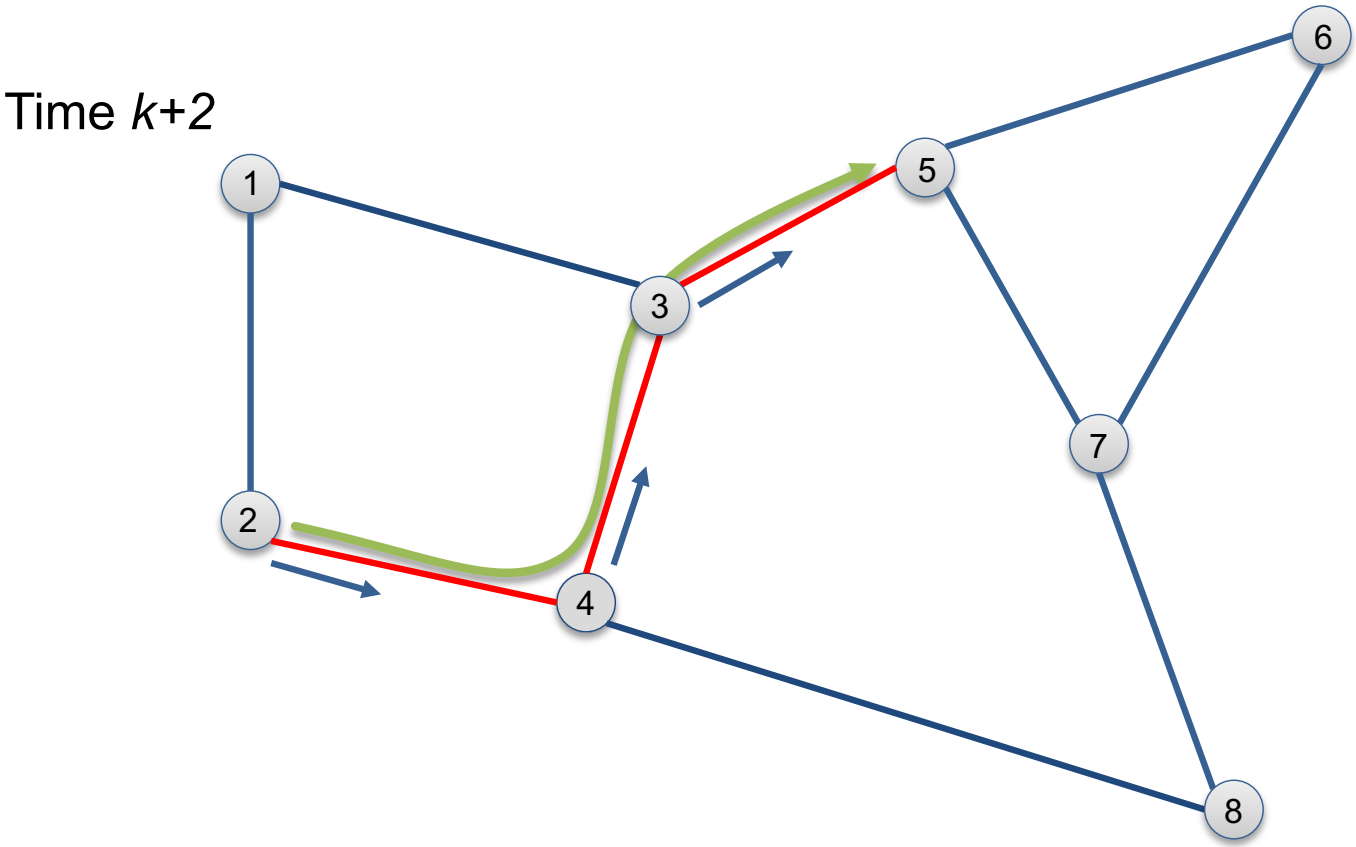

Convergence is reached when $x_1^K = x_2^K = \ldots = x_N^K$

Data Falsification by Compromised Nodes

Byzantine attacks inject false data into the network

Time *k*+2

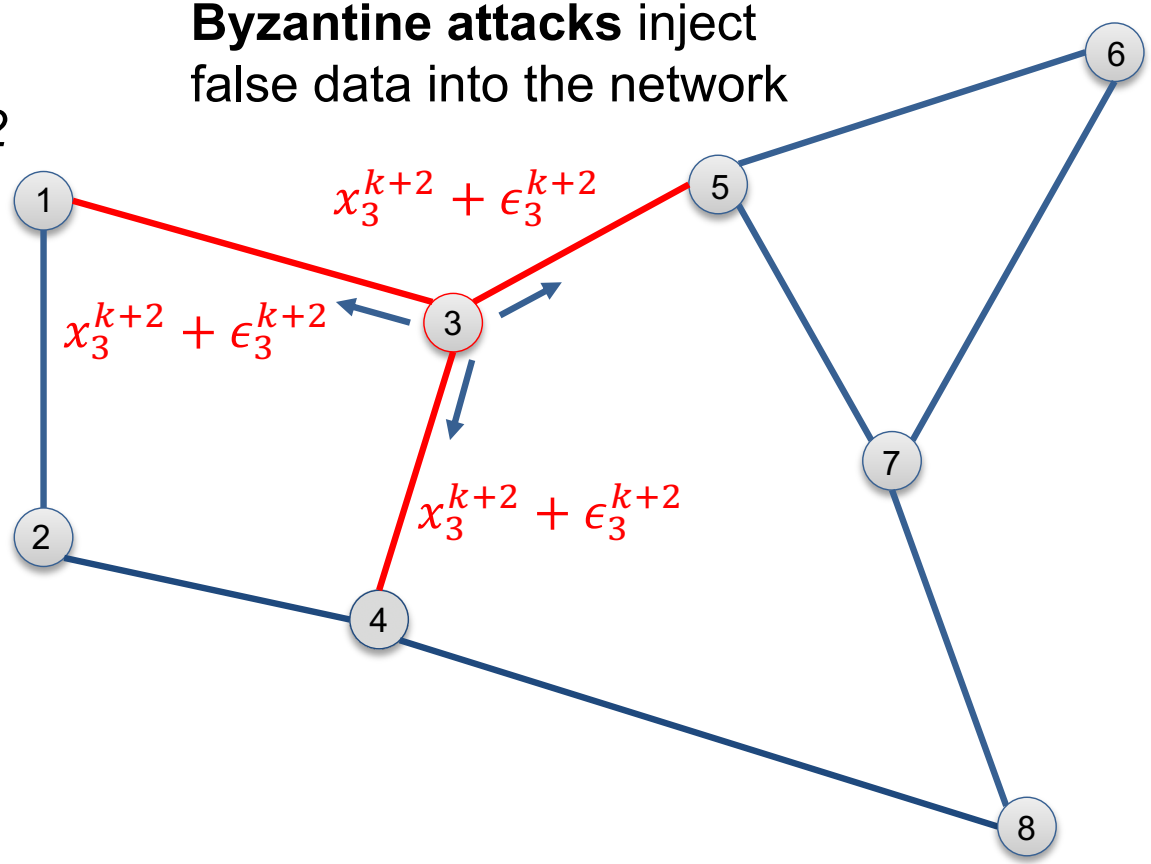

Why ns-3?

- Discrete-event network simulator intended for research use
- Free software licensed under GNU GPLv2
- Used in over 1000 peer-reviewed publications
- Publicly available models to extend functionality
- Compares favorably to other simulation tools
- Has been shown to scale with parallelization

Available at https://www.nsnam.org/

Computing Hardware

- My Computer
  - MacBook Air, Early 2015
  - 2 cores

- My Other Computer
  - CTS-1 Commodity Cluster
  - 96768 cores

More info: https://computation.llnl.gov/computers/commodity-clusters

Notional Scenario

- 1000 observer nodes
- 1 source node
- 10 km x 10 km x 3 km region of interest
- Observers set at elevations from 10 m to 3 km
- Source set on the ground

Centralized vs. Decentralized Performance

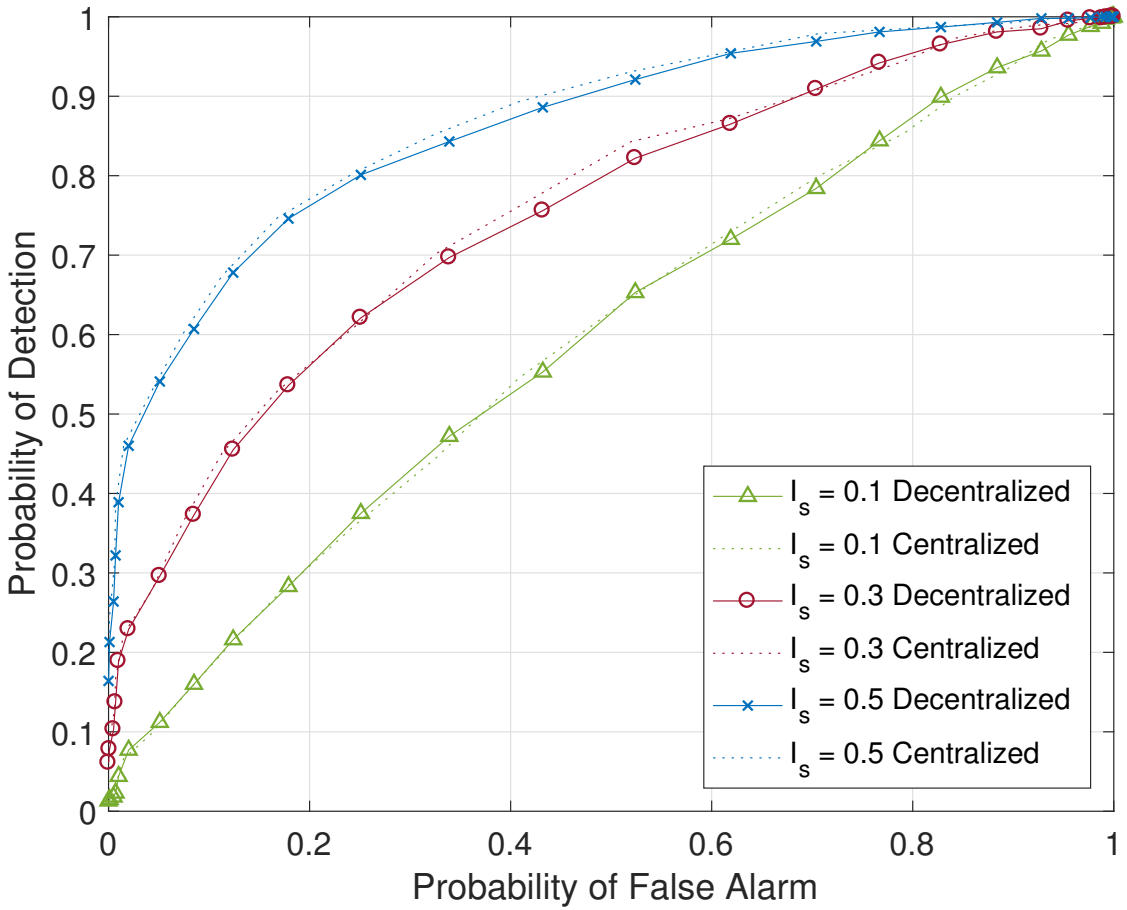

Convergence achieved within about 20-40 rounds

When You Add a Byzantine…

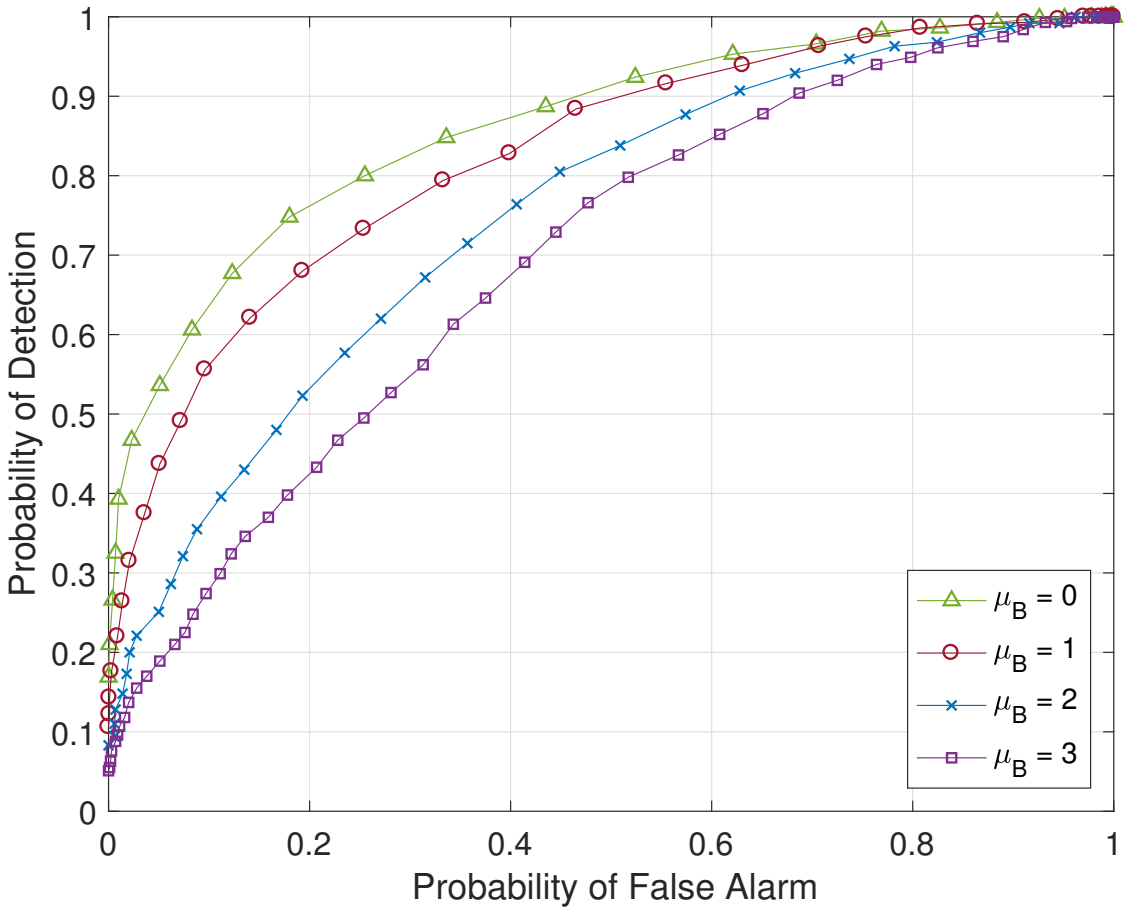

Only 1 Byzantine degrades a network of 1000 nodes

Robust ADMM Helps Mitigates the Attack

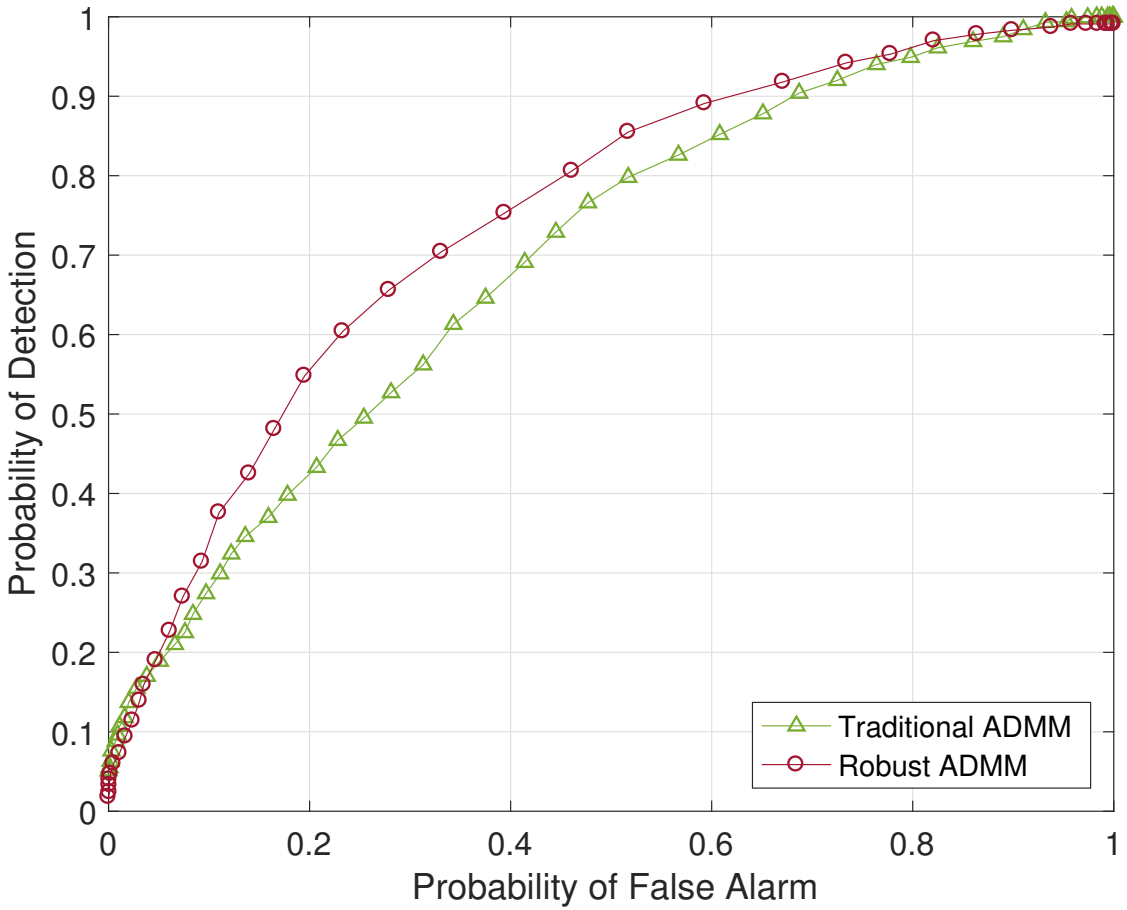

Greater impact on smaller networks

Small Performance Hit Without Byzantines

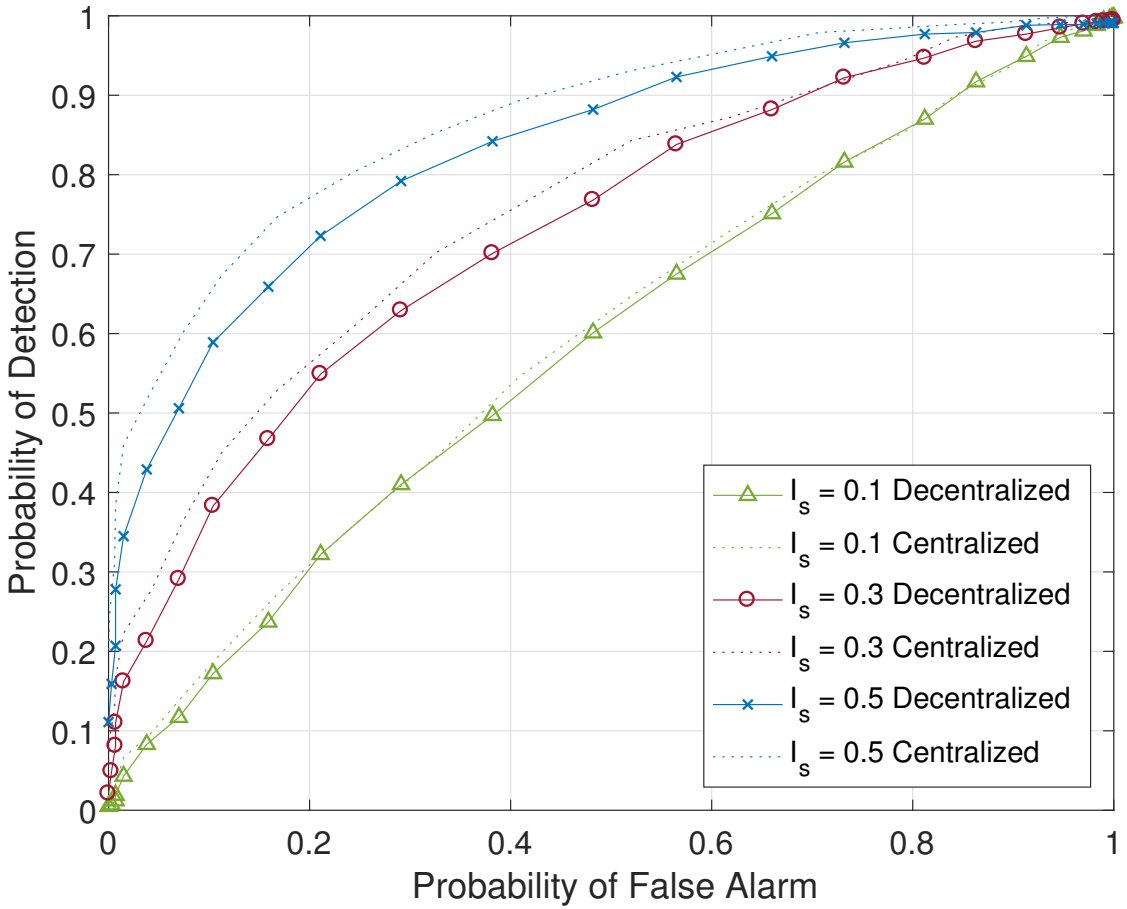

Better than dealing with a Byzantine, though!

Summary

- In a no-cyber world, decentralized performance is great
- Byzantine attacks can severely degrade performance
- Robust ADMM approach can help mitigate loss
- Small losses without Byzantines present
- Which would you rather deal with?

More questions? Email: yen6@llnl.gov

Lawrence Livermore
National Laboratory