



Cynthia Phillips
Sandia National Laboratory
Friday, March 31, 2017
3:00 PM
Lindley Hall, Rm. 102

Cooperative Computing for Autonomous Data Centers

Abstract: We present a distributed model for graph computations motivated by limited information sharing. Two autonomous entities have collected large social graphs. They wish to compute the result of running graph algorithms on the entire set of relationships. Because the information is sensitive or economically valuable, they do not wish to simply combine the information in a single location and then run standard serial graph algorithms. We consider two models for computing the solution to graph algorithms in this setting: 1) limited-sharing: the two entities can share only a polylogarithmic size subgraph; 2) low-trust: the two entities must not reveal any information beyond the query answer, assuming they are both honest but curious. We believe this model captures realistic constraints on cooperating autonomous data centers. We present results for both models for s-t connectivity. In the limited-sharing model, our results exploit social network structure to exchange $O(\log^2 n)$ bits, overcoming polynomial lower bounds for general graphs. In the low-trust model, our algorithm requires no cryptographic assumptions for 3 or more centers and does not even reveal node names. We then consider low-communication algorithms in this setting for the planted clique problem. We conjectured a sociologically-justified structural property of social networks with only human nodes. Given this property, we developed an $O(\log^2 n)$ -communication algorithm for a planted clique with $\log n$ randomly-selected nodes and adversarial edge distribution for 2 centers. However, validation experiments on real networks failed because we needed a more nuanced social network model and because these real networks contain significant non-human behavior. We describe a method for cleaning some non-human behavior given only topology. We'll give some evidence the notion is correct and potentially important. The cleaning gives more human networks for algorithm experimentation and indicates properties that might be relevant for future algorithm research.

This is joint work with Jon Berry (Sandia National Laboratories), Michael Collins (Christopher Newport University), Aaron Kearns (University of New Mexico), Jared Saia (University of New Mexico), and Randy Smith (Sandia National Laboratories).

Biography: Cynthia Phillips is a senior scientist at Sandia National Laboratories. She received a B.A. in applied mathematics from Harvard University and a PhD in computer science from MIT. She has historically worked in combinatorial optimization, algorithm design and analysis, and parallel computation with elements of operations research. Her work has spanned theory, general solver development, and applications. Most recently, she has worked in "big data" areas, including streaming, write-optimized data structures, complex (social) network analysis, trajectory analysis and I co-design of algorithms and architectures for extreme-scale future machines. Sandia algorithms researchers work in diverse areas. Her current or recent research areas include scheduling, network and infrastructure surety, sensor placement, wireless networks, integer programming, graph algorithms, and cybersecurity. In the more distant past, she has worked on computational biology, experimental algorithmics, and quantum computing.

