Network Mapping and Measurement Conference 2012

Arizona State University
March 19–20, 2012

Program for Monday, March 19, 2012

8:30-9:30am Breakfast in Artisan Court at the Brickyard (BYAC) Lobby

9:30-10:30am Invited Talk, "Compressive Demodulation of Mutually Interfering Signals," (abstract page 3)
Robert Calderbank, Duke University.

10:30-11:00am Morning Break

11:00-11:40am "Sublinear Time Recovery of Sparse Signals," (abstract page 3)
Charles J. Colbourn, Arizona State University.

11:40-1:30pm Lunch

1:30-2:10pm "An Introduction to the Challenge Problems," (abstract page 3)
John Triechler, Raytheon Applied Signal Technology, Inc.

2:10-2:50pm "Sensor Provisioning for Multistatic Tracking of Target Reflectors," (abstract page 5)
Alfred O. Hero III and Greg Newstadt, The University of Michigan.

2:50-3:20pm Afternoon Break

3:20-4:00pm "Fast Packet-Free Inference of Wireless Topologies: A Control-Theoretic View of Networks," (abstract page 5)
Sean Warnick and Daniel Zappala, Brigham Young University.

4:00-4:40pm "Security and Discoverability of Spread and Flow Dynamics: A Network Control Theory Perspective," (abstract page 6)
Mengran Xue and **Sandip Roy**, Washington State University;
Sean Warnick and Anurag Rai, Brigham Young University.

4:40-5:00pm Planning meeting for NMMC 2013.

Both the BYAC and BYENG buildings have WiFi access throughout. To use it, select the SSID "asu guest" and accept ASU's policy of use.

Network Mapping and Measurement Conference 2012

Arizona State University
March 19–20, 2012

Program for Tuesday, March 20, 2012

8:30-9:30am Breakfast in Artisan Court at the Brickyard (BYAC) Lobby

9:30-10:30am Invited Talk, "Radio Tomography: Environmental Inference from Wireless Network Signal Strength Measurements," (abstract page 7)
Neal Patwari, The University of Utah.

10:30-11:00am Morning Break

11:00-11:40am "Internet Topology Inference and the High Rank Matrix Completion Problem," (abstract page 8)
Brian Eriksson, Boston University.

11:40-12:20pm "Socioscope: Spatio-Temporal Signal Recovery from Social Media," (abstract page 8)
Robert Nowak, The University of Wisconsin.

12:20-2:00pm Lunch

2:00-2:40pm "State-space Models for Dynamic Networks," (abstract page 8)
Kevin S. Xu and Alfred O. Hero III, The University of Michigan.

2:40-3:20pm "Rate Control for Wireless Networks: How Good are the Models?," (abstract page 9)
Daniel Zappala and Sean Warnick, Brigham Young University.

3:20-4:00pm "Group Gossip: Distributed Consensus Through Multilateral Wireless Exchanges," (abstract page 9)
Waheed U. Bajwa, Rutgers University.

Both the BYAC and BYENG buildings have WiFi access throughout. To use it, select the SSID "asu guest" and accept ASU's policy of use.

Abstracts (in talk order)

Compressive Demodulation of Mutually Interfering Signals

Robert Calderbank, Duke University

The challenge of Multiuser Detection (MUD) is that of demodulating mutually interfering signals, given that at any time instant, the number of active users is typically small. The promise of compressed sensing is demodulation of sparse superpositions of signature waveforms from very few measurements. This talk will describe how methods of compressed sensing can be applied to Multiuser Detection and how the results can be applied to communication in wireless sensor networks.

Biography: Robert Calderbank is Dean of Natural Sciences and Professor of Electrical Engineering at Duke University where he directs a research program at the interface of signal processing and wireless communication. Prior to joining Duke in 2010, he was Director of the Program in Applied and Computational Mathematics at Princeton University, and before joining Princeton in 2004 he was Vice President for Research at AT&T. At the start of his career at Bell Labs Professor Calderbank developed technology that was incorporated in a progression of voiceband modem standards that moved communications practice close to the Shannon limit. Together colleagues at AT&T Labs he showed that good quantum error correcting codes exist and developed the group theoretic framework for quantum error correction. He is a co- inventor of space-time codes for wireless communication, where correlation of signals across different transmit antennas is the key to reliable transmission. Professor Calderbank was elected to the US National Academy of Engineering in 2005.

Sublinear Time Recovery of Sparse Signals

Charles J. Colbourn, Arizona State University

Sparsity underlies not only the ability to compress a signal with many coordinates, but the ability to compress while sensing or sampling. A wide variety of techniques have been developed to undertake such compressive sensing, under different assumptions about the structure of the signal, its sparsity, the amount and type of noise, and the recovery technique. Often techniques that exhibit better compression necessitate more complex recovery methods, and hence are limited to relatively small signals. We have developed an hierarchical technique for compressive sensing that uses many small measurement matrices, each potentially having its own recovery method, into one larger measurement matrix. In this setting, recovery for the larger signals can be shown to be easily reducible to recovery of many projections onto the (much smaller) ingredient matrices. This reduction underlies a general method for obtaining recovery algorithms whose running time is sublinear in the length of the signal, while permitting excellent compression. The methods developed rest on a combinatorial framework for constructing measurement matrices using hash families, which in turn can be constructed from well-known algebraic structures.

This is joint work with Violet R. Syrotiuk (Arizona State University) and Daniel Horsley (Monash University).

An Introduction to the Challenge Problems

John Triechler, Raytheon Applied Signal Technology, Inc.

Problem 1A: Developing a Control Theoretic Representation of Network Behavior

An elaborate theory of measurement and control has been developed over the past sixty years for linear electrical and mechanical systems, and its value has been proven by its broad and successful practical application. An important attribute of this theory is its ability to explain the behavior of complex systems and to provide simple mathematical tests which can determine the ability of a system operator to measure its performance and control its behavior. The last twenty years has seen the parallel growth of communications networks of ever growing complexity. The intent of this challenge problem is to examine the degree to which a theory of observability and controllability can be developed for these ever evolving and ever more complex communications networks. If so, then there would clearly be new opportunities for their efficient management using closed-loop control systems.

To develop this theory we desire to find suitable network-centric definitions for: (a) "Identification," the determination of the topology of a network; (b) "Observability," the determination of traffic flows through the network; and (c) "Controllability," the control of traffic flows through the network. While the terms used here are taken from linear systems control theory, and the concepts are roughly the same, there is certainly no assumption that the same mathematical constructs will apply. Thus the challenge is to develop an appropriate set of mathematical definitions which also have a reasonable intuitive meaning to engineers and practitioners.

Problem 1B: Using Control Theoretic Concepts in the Limit of Restricted Measurements

In very few networks is it possible to make all the measurements that one would prefer, or even require, to have in order to assess the ability to identify, observe, and control a network. This portion of the challenge problem is to address just that. Specifically, the question is, how measurements, and of what type, are required to identify the topology of a network, to determine its switching or routing behavior, and ultimately to control (manage) its flows.

In order to guide work on this challenge, it should be addressed within the following framework: (a) Activity measurements in the network will be made on the links, not the nodes. (b) Timing accuracy of the activity measurements will be precise enough to judge simultaneity, but not to determine the sequence of links over which a call, packet or message passes. (c) Transit times on the links are assumed to be instantaneous (for now), thus removing propagation delay as a complicating issue. (d) It should be assumed (for now) that no packets, messages, or calls are lost, and that the probability of detection on a link is 100% and that there are no false alarms. (e) It is not possible to measure activity on all links. An important implication of this last assumption is that it is quite possible that the network under examination is not "observable," by the definition developed in Problem 1A. A key aspect of this challenge problem is to determine the effect of this nonobservability has on the other two of the attributes of identifiabiity and controllability.

Problem 2: Optimal Methods for Multistatic Tracking of Target Reflectors

The availability of high speed signal process processing and sophisticated signal processing algorithms has made it possible to move beyond traditional radio ranging systems (radars), where a single transmitter and single receiver are co-located, and on to systems with multiple transmitters and multiple receivers, all of which are geographically dispersed. This challenge problem is to determine how to obtain the best performance from such a system.

In order to constrain the problem, we make the following assumptions: (a) The objective is to locate and then track $N \geq 1$ potentially moving reflectors operating in field of view of the transmitters ("illuminators") and the receivers. (b) There are L, where $L \geq 1$, illuminators whose static locations and technical characteristics are known, but which cannot be modified, either statically or in real time. (c) There are M, when $M \geq 1$, receivers whose locations are known, whose processing can be changed at will, and who can communicate freely with each other.

There are several important variations on this problem which should be considered:

- The potential for moving the receivers, or, even the use of moving receivers.
- The constraints on performance imposed by the radio propagation behavior between each pair of illuminators and receivers. To constrain this a bit, consider two limiting cases:
  - HF over-the-horizon radar using linear FM to obtain good range resolution but propagating through time-varying, dispersive media, with the potential of needing to locate and track reflectors over an area of tens of thousands of square km, and
  - Wideband CDMA using 5 MHz-wide direct-spread signals operating in the UHF band with very little atmospheric dispersion, with the potential of locating and tracking reflectors over an area of a few cells in a wireless system (nominally 10 square km).
- The advantages, if any, of using illuminators operating at substantially different frequencies.

The figure of merit for this challenge problem should be some function of the number of reflectors within the combined field of view which can be unambiguously detected and then tracked, and the spatial accuracy with which they can be tracked.

Sensor Provisioning for Multistatic Tracking of Target Reflectors

Alfred O. Hero III and Greg Newstadt, The University of Michigan

Sensor provisioning is the problem of determining the sensing/computational resources required to accomplish a complicated system level task, e.g. tracking or discriminating between N targets. It is a central problem in both traditional and non-traditional radar systems, where the number of targets can easily exceed the number of radars. Here we present a method for conservative sensor provisioning that guarantees a prescribed level of system performance, e.g. multiple target detection and position uncertainty levels, regardless of the scenario. We provide fundamental performance limits, such as the number of targets that can be tracked with bounded uncertainty, the system occupancy rate, and the load margin, e.g., the amount of resources available for other tasks. We provide multiple operational contexts, including provisioning for (a) passive radar systems using CDMA signals from multiple fixed transmitters and receivers and (b) multistatic synthetic aperture radar (SAR) systems where the locations of the receivers are allowed to move.

Fast Packet-Free Inference of Wireless Topologies: A Control-Theoretic View of Networks

Sean Warnick and Daniel Zappala, Brigham Young University

Interconnected dynamic systems are a pervasive component of our modern infrastructures. The complexity of such systems can be staggering, which motivates simplified representations for their manipulation and analysis.

Historically, mathematical representations for such complex phenomena have focused on either input-output representations or state space models. Although both types of representation capture the dynamics of such systems, they describe very different perspectives of the structure of a system: in terms of structure, state space models describe a rich interconnection structure of the system, while input-output models only indicate the existence of some path between each input and output. Since these representations both equally capture system dynamics, yet the state space model also describes much more structural information, it is not surprising that identifying the state space model from data requires knowing additional information about the system.

This additional information required to specify the state space realization of a system can be impossible to acquire in practice. Essentially it is equivalent to knowing the basis, or representation and units of all system variables, including those that you do not directly observe as system outputs. Developing a theory for the identification of system structure thus demands a new system representation, one that describes more structure than the input-output description but less than the state space realization (and thus also demands less additional information for identification from data).

Dynamical structure functions offer such a representation. They describe the causal relationships among manifest variables (inputs and outputs) but offer no insight about the structure of other, hidden system variables. As a result, their identification requires less additional information than identifying the full state space representation. In particular, a systems dynamical structure function can be identified in $O(p)$ experiments, where p is the number of measured outputs. Each experiment involves measuring all outputs in response to a particular excitation.

This work applies this new theory of network identification (which addresses Design Challenge 1A) to that of discovering the "topology" of a wireless mesh network (which addresses Design Challenge 1B). In this context, topology refers to directed contention and interference relationships between links.

To accomplish this, we present a model of the contention dynamics resulting from the carrier sensing behavior of 802.11. We offer some preliminary validation for the model from NS3 simulations and experimental results reported in the literature. This model then provides a state space description of the system dynamics; coupled with a particular network topology, this representation provides ground truth from which we test our identification methods.

A network discovery algorithm is then presented. The algorithm perturbs the rates of each link away from its current equilibrium, one at a time, and then measures the resulting realized rates on all network links. This information yields a graphical representation of the systems network structure near the corresponding equilibrium; as network traffic demands change, the procedure can be repeated to offer a new view of the active structure. One important feature of this procedure is its ability to provide information similar to measuring the network interference map, but with only $O(p)$ experiments (instead of $O(p^2)$) and without shutting the network down while the system is calibrated.

Another interesting point is that the procedure is packet-free, in the sense that only link rates are measured, not packet delays, etc. As a result, the discovered topology is not necessarily a map of pathways packets can travel on the network. To drive this point home, we consider a competitive situation where the private wireless communication among intruders disrupts the performance of an existing network. Although the topology of the existing network has not changed, in the sense of how packets may flow, recalibration of its dynamical structure function reveals the presence of the intruders as a change in the coupling structure among its links.

Security and Discoverability of Spread and Flow Dynamics: A Network Control Theory Perspective

Mengran Xue and Sandip Roy, Washington State University

Sean Warnick and Anurag Rai, Brigham Young University

The operation of infrastructure networks (e.g., transportation, power, and communication networks) is becoming increasingly complex: these networks are increasingly forced to operate near their limits, while being subject to an ever-wider family of threats and uncertainties; at the same time, new cyber- systems are being intertwined with these infrastructures to assist with decision-making and automation, further complicating their dynamics. As these networks become increasingly complex, individual stakeholders whether network planners, adversaries, or users are increasingly limited in their knowledge of the system. These stakeholders may not have full knowledge of networks components and topology, and they almost certainly will be able to measure only a small part of the networks dynamics in real time. Given these limits in knowledge and information flow, the stakeholders increasingly need to estimate critical information about network structure and dynamics from highly limited measurements.

Motivated by this need for estimation in complex networks, our group has pursued research in two directions: 1) development of foundational tools based on control-theory concepts for estimation of network structure and dynamics; and 2) definition/analysis of security and, conversely, discoverability measures as indications of how well a stakeholder can estimate dynamics/structure. The foundational tools that we have developed have been concerned with state and mode (eigenvalue) estimation problems for several canonical network models, including models for synchronization, infection spread, and autonomous-vehicle-team motion [1-3]. These state and mode estimation problems can at their essence be addressed using classical filtering and smoothing algorithms from the control-theory and signal-processing literature. What is novel in our efforts is that we tie the estimators structure and performance to the topology of the network and the observation locations; that is, we examine what characteristics of the network and the observation locations either facilitate or hinder estimation, so as to inform network design and measurement placement. Using these results, we have also pursued characterization/design of security in these canonical network models, where security is defined in terms of an adversarys performance in estimating the network dynamics from local measurements.

We believe that the control-theoretic network-estimation tools that we have developed are promising for allowing estimation of dynamics and identification of topological structure in information-transmission networks (which addresses Challenge Problem 1B). However, the network-estimation theory developed in [1-3] needs to be advanced in several ways to develop practical tools for communication (information-transmission) networks. Three critical needs:

1. The network-theoretic characterization of estimation needs to be developed for a broader class of canonical models, including especially flow- and queueing- network models (which can represent information-transmission processes). The canonical spread and synchronization models for which network-theoretic estimation has been achieved are relevant to information-transmission-related applications (for instance, to represent gossip algorithms for information flooding, spread of a computer virus, or a consensus algorithm), however many communications processes are better represented by flow- and queueing- network models, and are open models (in the sense that traffic flows into and out of the model). We need to enhance our network-theoretic estimation studies to incorporate such dynamics.

2. Thus far, our studies of network estimation have assumed that the network dynamics is deterministic, and only measurements are subject to noise. In fact, communication network dynamics often require stochastic models: routing of packets, traffic inflows, and perhaps even network topologies are best represented as stochastic

processes. We will pursue network-theoretic estimation in models that abstractly capture some of these uncertainties.

3. In our network-estimation efforts thus far, we have assumed that the states of network components (nodes) are measured. In fact, as indicated in the challenge problem, traffic flows on links, rather than nodal quantities, are likely to be measured in information-flow networks. Thus, the network-theoretic estimation results need to be adapted to the case where measurements are concerned with link quantities. Additionally, if estimation of link characteristics is of interest, the canonical network models may need to be modified to directly track link traffic rather node characteristics.

In this presentation, we will provide a tutorial on mode and state estimation in canonical network models, focusing especially on network-theoretic characterization of estimator performance in a canonical model for spread. We will then present preliminary explorations in the three directions noted above, as a step toward understanding state and topology inference in information-transmission networks. Numerical results and simulations will be included throughout the presentation to illustrate the technical development.

References:

1. M. Xue, A. Rai, E. Yeung, S. Roy, Y. Wan, and S. Warnick, Initial-Condition Estimation in Network Synchronization Processes: Graphical Characterizations of Estimator Structure and Performance, in Proceedings of AIAA Guidance, Navigation and Control Conference, Portland, OR, August 2011.

2. S. Roy, M. Xue, and S. K. Das, Security and Discoverability of Spread Dynamics in Cyber-Physical Networks, to appear in IEEE Transactions on Parallel and Distributed Systems, 2012.

3. M. Xue, W.Wang, and S. Roy, Security Concepts for the Dynamics of Communicating Autonomous-Unmanned-Vehicle Networks, submitted to Automatica.

Radio Tomography: Environmental Inference from Wireless Network Signal Strength Measurements

Neal Patwari, The University of Utah

The received power on a link between two static wireless devices is changed when a person stands near the line those two devices. We describe multiple methods to infer a person's location in a room or building using received signal strength (RSS) changes measured in a static wireless network. We describe tomographic imaging approaches and statistical inversion methods to this problem, which we call device free localization (DFL) because the person does not participate in the system by carrying any device. We show that a person can be tracked with average errors less than 1 meter, even when using a network outside of a building to perform through-wall imaging and tracking. On some links, RSS changes even when a person is stationary, due to their inhalation and exhalation. We describe algorithms to use these changes to reliably and accurately estimate breathing rate. In summary, we describe how networks made of standard wireless devices might be used to learn characteristics of the changing environment in which they are deployed.

Biography: Neal Patwari received the B.S. and M.S. degrees from Virginia Tech, in 1997 and 1999, respectively, and the Ph.D. degree from the University of Michigan, Ann Arbor, in 2005, all in electrical engineering. He was a research engineer with Motorola Laboratories, FL, between 1999 and 2001. In 2006, Prof. Patwari joined the Electrical and Computer Engineering Department at the University of Utah in Salt Lake City where he is an Assistant Professor with an adjunct appointment in the School of Computing. He directs the Sensing and Processing Across Networks (SPAN) Laboratory, which performs research at the intersection of statistical signal processing and wireless networking. Dr. Patwari received the NSF CAREER Award in 2008, the 2009 IEEE Signal Processing Society Best Magazine Paper Award, and the 2011 University of Utah Early Career Teaching Award. He has served on technical program committees for IEEE conferences SECON, ICDCS, DCOSS, ICC, RTAS, WoWMoM, ICCCN, and MILCOM. He is an Associate Editor of the IEEE Transactions on Mobile Computing.

Internet Topology Inference and the High Rank Matrix Completion Problem

Brian Eriksson, Boston University

The ability to recover Internet router-level connectivity is of importance to network managers, network operators and the area of security. As a complement to the heavy network load of standard active probing methods, which scale poorly for Internet-scale networks, recent research has focused on the ability to recover Internet connectivity from passively observed Internet traffic. Infrastructures exist that record hop distances from N end host computers to a set of n monitoring points throughout the Internet. Due to the nature of these observations, these $[n \times N]$ hop count matrices will be massively incomplete. This poses a matrix completion problem, with the incomplete distance matrix being potentially full-rank in this application. Fortunately, computers tend to be clustered within subnets having a small number of access points to the Internet at large, therefore the columns of the matrix lie in the union of low-rank subspaces.

This talk considers the problem of completing a matrix with many missing entries under the assumption that the columns of the matrix belong to a union of multiple low-rank subspaces. This generalizes the standard low-rank matrix completion problem to situations in which the matrix rank can be quite high or even full rank. For an $[n \times N]$ matrix whose columns lie in a union of at most k subspaces, each of rank $< r$, we show that under mild assumptions each column of the matrix can be perfectly recovered with high probability from an incomplete version so long as at least $rN \log^2(n)$ entries are observed uniformly at random.

Socioscope: Spatio-Temporal Signal Recovery from Social Media

Robert Nowak, The University of Wisconsin

Social media sources are rich in information, but the sources are decentralized, unreliable, and uncalibrated. This leads to unique and challenging data fusion problems. This talk describes Socioscope, an inference algorithm specifically designed to reconstruct spatiotemporal maps of events of interest from social media data. The Socioscope is based on two types of observational data collected from social media sources. The data are counts of the occurrences of specific types of messages/reports. We will refer to these occurrences as "events." One set of data are "target events" related to specific topics of interest. The other data are "miscellaneous events" that are used to calibrate the number of observed target events relative to the background level of the sources. These events are extracted from social media streams using natural language processing (NLP) filters. Each event is marked with a location and time stamp, although in general these may be unreliable or missing, issues we will deal with in our framework. The event data are modeled as realizations of a marked Poisson process. The problem of recovering the underlying spatiotemporal distribution of target events is posed as a Poisson inverse problem. A penalized maximum likelihood estimator is proposed and experiments with real Twitter data demonstrate its performance.

State-space Models for Dynamic Networks

Kevin S. Xu and Alfred O. Hero III, The University of Michigan

Networks are ubiquitous in science, serving as a natural representation for many complex physical, biological, and social phenomena of interest. Significant efforts have gone into the development of statistical models for networks. Most existing work has focused on modeling static networks, which represent either a single time snapshot or an aggregate view over time of the phenomenon of interest. A static network representation is in many cases an oversimplification that is unable to capture important features of the phenomenon such as its temporal dynamics. The natural extension is to consider dynamic, time-evolving networks.

To date, there has not been much work on statistical modeling of dynamic networks. We propose a state-space model for discrete-time dynamic networks that combines two types of statistical models: a static model for individual network snapshots and a temporal model for the state evolution over time. The network snapshots are modeled using the stochastic block model, a simple parametric model for static networks dating back to the 1960s, while the temporal dynamics are modeled by a standard linear dynamic system. We fit the model using a modified Kalman filter augmented with a local search strategy. We demonstrate the proposed method for state tracking and link prediction in dynamic networks.

Rate Control for Wireless Networks: How Good are the Models?

Daniel Zappala and Sean Warnick, Brigham Young University

Developing an accurate model of wireless networks would enable improved applications in the areas of performance, monitoring, and security. It is difficult to know how well a wireless network should operate, when all transmissions must share airspace with other wireless devices. Likewise, when performance is not meeting expectations, it is difficult to know whether this is due to capacity constraints, interference, or misbehaving or malicious devices. Addressing these issues requires a combination of modeling, optimization, implementation, and experimentation. Our goal is to unify these various approaches to create optimal yet practical protocols for wireless mesh networks.

One focus of our current research is to build a rate control protocol that can optimally allocate rates to the set of active flows in a wireless mesh network. Rate allocation in wireless networks has been formulated as a network utility maximization problem by numerous papers, where the objective function is typically a sum of utility functions for each flow's sending rate. A set of constraints is used to model the unique characteristics of the wireless network, such as carrier sensing or interference constraints. The solution to this optimization problem will then yield a set of rates that maximize network utility. When the optimization problem is convex, it can often be translated into a distributed rate control algorithm, making it practical to deploy in a wireless network.

Relatively few of these papers have attempted to implement their solution, providing little insight into whether these modeling efforts have succeeded in capturing the complexities of carrier sensing and interference in wireless networks. We have developed new models that take into account interference that is prevalent in wireless networks and shown through numerical simulations that these new models provide greater accuracy and thus better network utility. Absent any experimental results, however, the question remains: how good are these models?

In this talk, we will describe our current progress toward answering the question of how good these models are in practical, daily use. To address this question we have been developing an open-source toolkit that makes it easier to build experimental transport protocols in user space. We will share our results from both experimentation in a mesh testbed and packet-level simulation. In particular, we will show how a fundamental assumption on how interference behaves is incorrect. Because this assumption underlies the modeling in this area, this has a major effect on whether it is better to ignore or accurately model interference.

Group Gossip: Distributed Consensus Through Multilateral Wireless Exchanges

Waheed U. Bajwa, Rutgers University

Distributed consensus is a moniker used to describe a set of problems in which the network nodes iterate to collaboratively compute a function of initial node values. Gossip algorithms, which rely on pairwise exchanges among the network nodes, constitute a class of techniques designed to address the linear consensus problem. In the case of wireless networks, many such algorithms fail to capture the broadcast and superposition nature of the wireless medium. In this talk, we present a more flexible family of linear consensus techniques that is constructed specifically for wireless networks. This family of algorithms, which we term group gossip, uses multilateral exchanges among the network nodes and relies only on local information and synchronization. Group gossip is not only superior to traditional gossip in terms of implementation over wireless networks, but it also permits a tighter bound on the convergence rate. In general, the upper bound on the convergence time is at most one-third compared to that for randomized gossip.

This work is in collaboration with Matthew Nokleby and Behnaam Aazhang (Rice University) and Robert Calderbank (Duke University).