Research Overview

Alejandro Ribeiro
aribeiro@seas.upenn.edu
Dept. of Electrical and Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104

Table of Contents

Contents

Research Overview

Although networks cannot be claimed to be a recent invention it is safe to assert that the irruption of networks into everyday life is a determinant feature of the late 20th century. Technological networks have transformed the way in which we interact with each other and the way in which we acquire and manipulate information. At the same time, we have come to realize that underlying networks of influence and interaction play a fundamental role in explaining the behavior of large scale natural and social phenomena. Understanding networks has thus emerged as one of the great intellectual challenges of the 21st century. My research is about the application of Signal and Information Processing tools to the design of networks, the development of network algorithms, and the understanding of networked phenomena.

For the most part my research is about the foundational understanding of these topics in what many have come to call the emerging field of Network Theory. There are myriad contexts in which networks play a fundamental role but a common thread that weaves them together is the local interaction versus global behavior dichotomy. In networked systems we only have knowledge and control on local actions but want to explain or enforce network-wide behavior. As a Network Theory researcher my expertise is in various tools that can be used to bridge this dichotomy. Besides this, my philosophy is to be agnostic with respect to tools and applications.

Agnosticism notwithstanding, the different projects that constitute my current research program accept a taxonomy that differentiates between situations in which the goal is to design algorithms suitable for implementation on networked infrastructures, analyze phenomena in which underlying network interactions explain a significant part of the observed behavior, or, design the network infrastructure itself:

Network algorithms. Algorithms are the typical deliverable of research efforts in Signal and Information Processing. E.g., the solution of an estimation problem is an algorithm to compute good estimates for a given data set. In network algorithms the information to be processed is distributed throughout the network and the goal is to design mechanisms to process this information based on local interactions between neighboring nodes. In, e.g., distributed estimation problems, observations are acquired by different nodes and the estimation algorithm intends to accumulate global information through message exchanges that are restricted to neighboring nodes. I currently have three active research projects in this space, namely, "Bayesian Network Games," "Distributed Optimization," and "Wireless Control Systems." These projects are discussed in sections C, E, and F, respectively.

Networked data structures. A networked data structure is one in which there is some notion of locality. In the count of pairwise message exchanges between members of a group most messages sent by a given individual are concentrated on a few recipients, the level of expression of a gene affects the level of expression of just a few other genes, the output of an economic sector has significant direct effect on just a few other sectors, and so on. Conventional information processing algorithms that ignore this local structure can be applied. However, if the networked structure is ignored we should not expect to be able to understand what is the role the network plays in explaining observed behaviors. In this problem category our goal is to perform analysis of networked data structures that do not ignore the fact that the network is, indeed, a network. The two projects that belong in this category are "Hierarchical Clustering in Asymmetric Networks" and "Authorship Attribution using Word Adjacency Networks." These projects are discussed in sections D and G, respectively

Network infrastructure. While it has taken humanity a century to realize, our society's infrastructure is about connectivity. The transportation infrastructure connects places, the energy infrastructure connects generators with consumers, the communication infrastructure connects people, and the market infrastructure connects savers with developers. In network infrastructure design problems we are given resources that we administer to establish desirable connectivity patterns. My work on network infrastructure is on the design of wireless communication networks. In such case we are given resources in the form of power budgets and bandwidth that we allocate to establish direct communication links in order to support end-to-end communication flows. On top of deciding which connections should be established among those possible for the given bandwidth and power, we also want to determine suitable routes, link shares, and admission control variables to support the required level of service. I am currently working on "Optimal Design of Wireless Communication and Networking Systems" and the design of "Communication Networks for Autonomous Robot Teams." These projects are discussed in sections A and B, respectively

A Optimal wireless communication and networking systems

In their classical work on limit distributions Gnedenko and Kolmogorov wrote that the "epistemological value of the theory of probability is based on this: that large-scale random phenomena in their collective action create strict, nonrandom regularity". In simpler words, randomness generates structure. It is often possible to infer properties of large-scale stochastic systems even if analogous deterministic counterparts are intractable. Randomness, in the form of fading – random variations in propagation coefficients between network nodes –, is inherent in wireless networks. My technical approach to the optimal design of wireless systems is to explore the structure introduced by fading so as to understand their fundamental properties. An example of the desired outcomes are recent results concerning the optimality of separating wireless networking problems in layers and per-fading state subproblems.

To explain these results consider a communication network deployed to support information flows with some required level of service. The goal of the communication network is to administer given resources to sustain said flows. In conventional wired networks the resource given is a set of physical connections between nodes; see Fig. 1. Supporting information flows requires finding routes between source and destination, determining link sharing strategies, and controlling the amount of traffic injected into the network. It was an early design specification to separate these problems in layers – routing, link and transport for the problems in the previous sentence – that operate more or less independently, and interact through standardized interfaces. While this was mostly a matter of ensuring inter-operability it is remarkable that this separation can be optimal. Specifically, it is possible to define separate per-layer optimization problems whose outcome coincides with the solution of a joint non-layered optimization. Mathematically, separability comes from the fact that the wired networking problem is convex – a linear program in fact. The Lagrangian dual problem can thus be solved instead. As it often happens, the Lagrangian exhibits a separable structure, which, as it turns out coincides with the conventional layers.

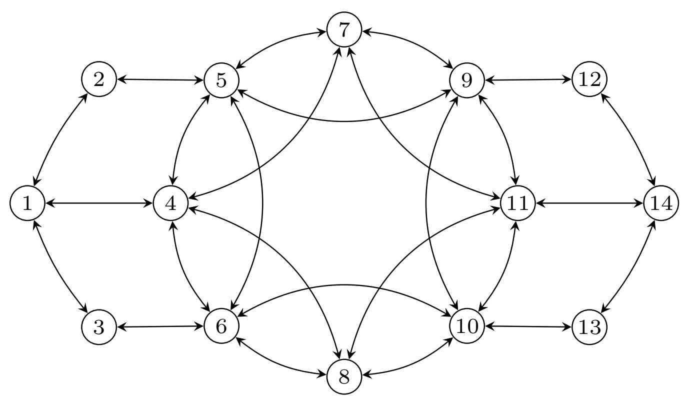

Figure 1: Wired network. In wired networks communication is possible between nodes sharing a physical connection. The goal of the communication network is to administer link resources to support information flows with some required level of service.

In a wireless network the given resources are not connections but bandwidth and power; see Fig. 2. Therefore, on top of routes, link shares, and rate control, a wireless networking problem entails determining which connections, among those possible for the given bandwidth and power, should be established to support the required level of service. Earlier approaches to wireless networking migrated the conventional layers and defined the power and frequency assignment as physical layer subproblems. This yields poor results though, and over time lead to the surge of cross-layer design as synonym of joint optimization across layers. Ultimately, the poor performance of layered wireless networks stems from the non-convexity of the joint cross-layer networking problem. As in wired networks, the Lagrangian exhibits a separable structure that can be mapped to layers. But non-convex problems have positive duality gap explaining the poor results of layered wireless networks.

In what constitutes an interesting example of structure introduced by randomness, we have proven that general wireless networking problems in the presence of fading, while non-convex, have zero Lagrangian duality gap. Exploiting the separability of the Lagrangian, this result yields the following principles:

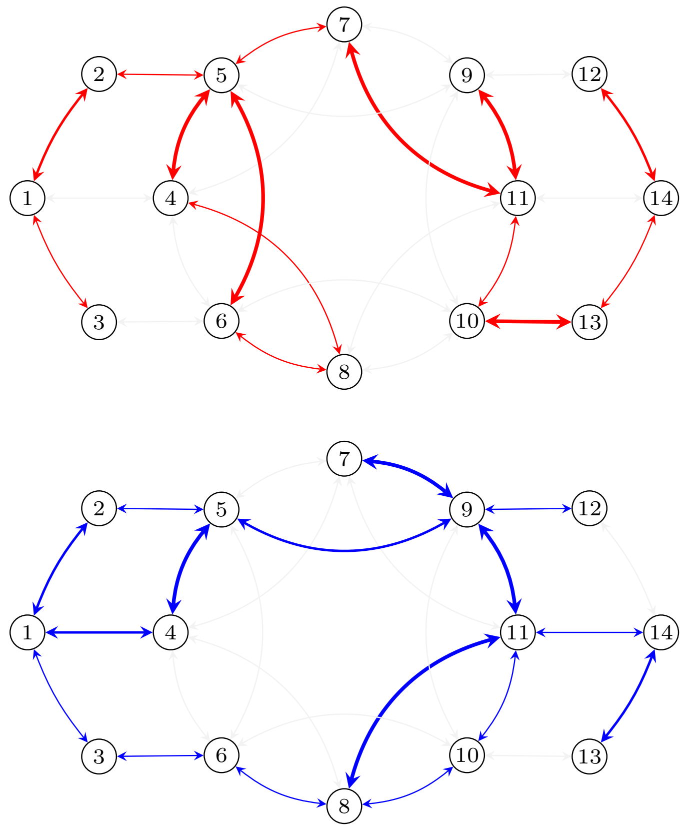

Figure 2: Wireless networks. In wireless networks (middle and bottom) links are established by a resource allocation decision to spend power and bandwidth on a specific connection for a specific fading state. Links can be different at different times and have time varying capacities.

First separation principle. This principle pertains to the separability of wireless networking problems into layers. It states that it is possible to define separate optimization problems to obtain optimal routes, optimal link capacity allocations, and optimal power/frequency assignments.

Second separation principle. Another difficulty in optimal wireless networking is the need to optimize jointly for all fading states. Given that fading coefficients take on a continuum of values, this is a variational problem that requires finding optimal functions of the fading coefficients. This principle states that network optimization is further separable in per-fading-state subproblems. The practical importance of this result is that it is not necessary to find optimal functions but only the values of the functions for those channels actually observed.

The separation principles hold under specific assumptions, e.g., networks operating in an ergodic setting and availability of perfect channel state information that restrict applicability of the separation principles to particular settings. Part of the research undertaken in the context of this project is to study the extent to which these assumptions can be lifted and what implications follow when this is not possible. Nonetheless, it has to be recognized that they do establish a fundamental property of wireless networks in the presence of fading that is not true for networks in a deterministic setting.

Our current research agenda utilizes the separation principles as motivation for part of the proposed research and as an example of the types of result that can be expected from the combination of optimization techniques and stochastic structure. Specific research activities include the construction of a generic framework to reduce the complexity of finding optimal operating points for wireless systems and the development of stochastic optimization algorithms to operate in environments where channel probability distributions are not known. We also apply this framework to the solution of specific problems in optimal wireless communications and networking. Of particular notice is the work on cognitive access algorithms and protocols. This is an effort to understand optimal wireless networking when terminals have different beliefs about the global state of the network.

A.1 Optimal wireless system design: Dual functions, duality gap, and separability

Operating variables of a wireless system can be separated in two types. Resource allocation variables $\mathbf{p}(\mathbf{h})$ determine instantaneous allocation of resources like frequencies and transmitted powers as a function of the fading coefficient $\mathbf{h}$. Average variables $\mathbf{x}$ capture system's performance over a large period of time and are related to instantaneous resource allocations via ergodic averages. A generic representation of the relationship between instantaneous and average variables is

$$\mathbf{x} \leq \mathbb{E}\left[\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)\right], \tag{1}$$

where $\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)$ is a vector function that maps channel $\mathbf{h}$ and resource allocation $\mathbf{p}(\mathbf{h})$ to instantaneous performance $\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)$. The system's design goal is to select resource allocations $\mathbf{p}(\mathbf{h})$ to maximize ergodic variables $\mathbf{x}$ in some sense.

An example of a relationship having the form in (1) is a code division multiple access channel in which case $\mathbf{h}$ denotes the vector of channel coefficients, $\mathbf{p}(\mathbf{h})$ the instantaneous transmitted power, $\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)$ the instantaneous communication rate determined by the signal to interference plus noise ratio, and $\mathbf{x}$ the ergodic rates determined by the expectation of the instantaneous rates. The design goal is to allocate instantaneous power $\mathbf{p}(\mathbf{h})$ subject to a power constraint so as to maximize a utility of the ergodic rate vector $\mathbf{x}$. This interplay of instantaneous actions to optimize long term performance is pervasive in wireless systems. A brief list of examples includes optimization of orthogonal frequency division multiplexing, cognitive radio, random access, communication with imperfect channel state information, and various flavors of wireless network optimization.

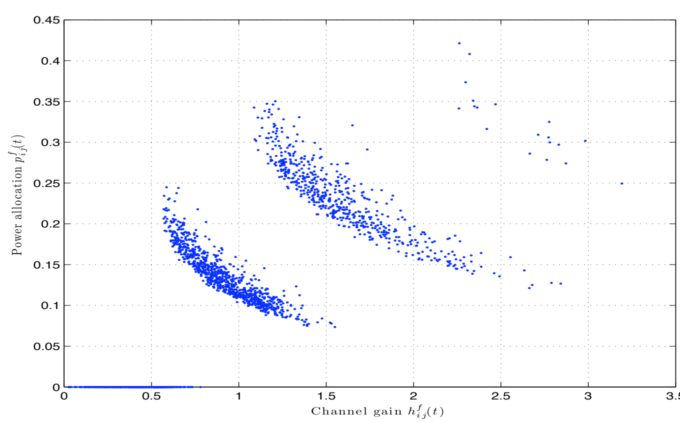

Figure 3: Power allocations in a wireless network using adaptive modulation and coding over an interference limited physical layer. Optimal power distribution consists of transitions between AMC modes and within each mode a cloud of allocations roughly proportional to the inverse channel.

In many cases of interest the functions $\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)$ are nonconcave and as a consequence finding the resource allocation distribution $\mathbf{p}^*(\mathbf{h})$ that maximizes $\mathbf{x}$ requires solution of a nonconvex optimization problem. This is further complicated by the fact that since fading channels $\mathbf{h}$ take on a continuum of values there is an infinite number of $\mathbf{p}^*(\mathbf{h})$ variables to be determined. A simple escape to this problem is to allow for time sharing in order to make the range of $\mathbb{E}\left[\mathbf{f}_1\big(\mathbf{h}, \mathbf{p}(\mathbf{h})\big)\right]$ convex and permit solution in the dual domain without loss of optimality. While the nonconcave function $\mathbf{f}_1(\mathbf{h}, \mathbf{p}(\mathbf{h}))$ still complicates matters, working in the dual domain makes solution, if not necessarily simple, at least substantially simpler. However, time sharing is not easy to implement in fading channels.

At the heart of this research endeavor there is a general methodology that can be used to solve optimal resource allocation problems in wireless communications and networking without resorting to time sharing. The fundamental observation is that the range of the expectation $\mathbb{E}\left[\mathbf{f}_1(\mathbf{h}, \mathbf{p}(\mathbf{h}))\right]$ is convex if the probability distribution of the channel $\mathbf{h}$ contains no points of positive probability. This observation can be leveraged to show lack of duality gap of general optimal resource allocation problems making primal and dual problems equivalent. The dual problem is simpler to solve and its solution can be used to recover primal variables with reduced computational complexity due to the inherently separable structure of the problem Lagrangians. The separation principles outlined above are a direct consequence of these observations but these properties can also be used to develop specific algorithms and protocols for optimal wireless communication and networking. An example of that is the ergodic stochastic optimization algorithm that we explain next.

A.2 Ergodic stochastic optimization algorithms

Wireless link capacities are expressed as expected values over fading channel realizations. Their evaluation, therefore, requires access to the channel probability distribution function (pdf). Assuming a particular fading model the channel pdf is acquired by estimating channel moments; e.g., for Raleigh fading estimating mean channels is sufficient. While this is a possible approach it does restrict practical applicability as fading models are only rough approximations of distributions observed in field deployments. To overcome this limitation we developed the ergodic stochastic optimization algorithm that learns channel distributions simultaneously with the determination of optimal operating points.

These algorithms are obtained through the implementation of stochastic subgradient descent in the dual domain and have been shown to exhibit almost sure convergence to optimal operating points in an ergodic sense. The ergodic stochastic optimization algorithm is implemented to find optimal operating variables for the network in Fig. 2 using adaptive modulation and coding over an interference limited physical layer. Nodes operate on 5 frequency bands, using direct sequence spread spectrum in each of them with spreading gain $S = 16$. Three AMC modes corresponding to capacities 1, 2 and 3 bits/s/Hz are used with transitions at SINR 1, 3 and 7. Fading channels are generated as i.i.d. Rayleigh with average powers 1/2 for the links $4 \leftrightarrow 7$, $5 \leftrightarrow 9$, $7 \leftrightarrow 11$, $9 \leftrightarrow 10$, $11 \leftrightarrow 8$, $10 \leftrightarrow 6$, $8 \leftrightarrow 4$ and $6 \leftrightarrow 5$

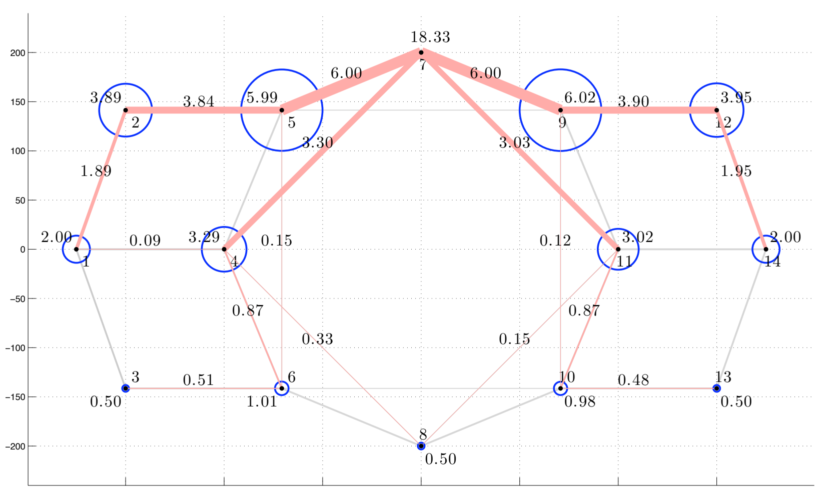

Figure 4: Optimal routes to node 7. Circles' areas are proportional to the amount of traffic terminals handle on behalf of terminal 7. Line widths are proportional to the packets routed through that link. Links of lower average quality are exploited, e.g., nodes 4 and 11 deliver information directly to node 7 bypassing the better quality links to nodes 5 and 9.

and 1 for the remaining links. Noise power is $N_i^f = 0.1$ for all terminals and frequency bands. The maximum average power consumption per terminal is $p_{\max} = 2$ – chosen so that if a terminal with 4 neighbors spreads power uniformly across all neighbors and frequencies the signal to noise ratio is 0dB. Powers p_i are also constrained to be positive, i.e., $p_{\min} = 0$. A spectral mask $\mathcal{P}(\mathbf{h}) := \{p_{ij}^f(\mathbf{h}) : 0 \leq p_{ij}^f(\mathbf{h}) \leq p_{\text{mask}}\}$ is further defined with $p_{\max} = 2$ the same value used to limit average power consumption. Link capacities and routing variables are constrained by $c_{\min} = r_{\min} = 0$ bits/s/Hz and $c_{\max} = r_{\max} = 6$ bits/s/Hz. Four flows with destination at terminals 1, 7, 8 and 14 are considered with all other terminals required to deliver at least $a_{\min} = 0.5$ bits/s/Hz and at most $a_{\max} = 2$ bits/s/Hz to each of these flows. Beyond that, the optimality criteria is sum rate maximization. Simulation results are presented in Figs. 3, 4, and 5.

Power allocation for the link from terminal 1 to terminal 4 is shown in Fig. 3. Optimal power distribution consists of transitions between AMC modes and within each mode a cloud of allocations roughly proportional to the inverse channel.

Routes to terminal 7 are sketched in Fig. 4. The circles' area is proportional to the amount of traffic that each terminal handles on behalf of the flow with destination at terminal 7. It is apparent that packets are indeed being delivered to the destination. It is also worth noting that links of lower average quality are exploited. E.g., nodes 4 and 11 deliver information directly to node 7 bypassing the better quality links to nodes 5 and 9.

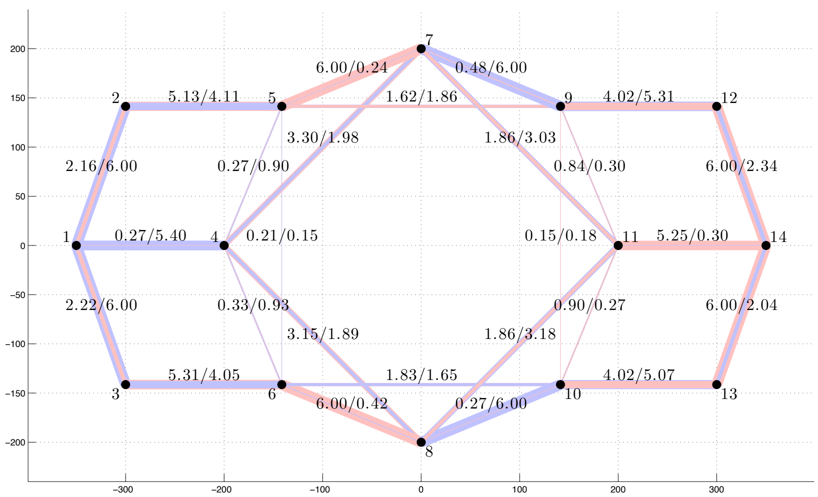

Figure 5: Optimal link capacities. The width of the link is proportional to the capacity of the link. For each pair of terminals $i < j$ the capacity c_{ij} between the smaller index i and the larger index j is shown first (in blue) and the capacity c_{ji} second (in red). Links to nodes 1, 7, 8 and 14 have the largest capacities as it should be of flows' destinations. Notice how some links, e.g. 11 to 9, despite having good average values are assigned small powers, thus having small capacity.

Ergodic link capacities are shown in Fig. 5. The width of the link is proportional to the capacity of the link. For each pair of terminals $i < j$ the capacity c_{ij} between the smaller index i and the larger index j is shown first (in blue) and the capacity c_{ji} second (in red). Links to nodes 1, 7, 8 and 14 have the largest capacities as it should be of flows' destinations. Also notice how some links, e.g. 11 to 9, despite having good average values are assigned small powers, thus having small capacity. This is a consequence of the joint – albeit separable – optimization of routes, link capacities and power allocations.

A.3 Acknowledgments and references

I started working on this project my last year at the University of Minnesota. Sincere thanks are therefore due to Prof. Georgios Giannakis. I also had the privilege of counting on the collaboration of Dr. Nikolaos Gatsis. After moving to the University of Pennsylvania this project morphed and expanded into the Ph. D. work of Yichuan Hu. The list of journal papers published in the context of this project is the following:

[1] Y. Hu and A. Ribeiro, "Optimal wireless communications with imperfect channel state information," *IEEE Trans. Signal Process.*, vol. (revised), January 2013.

[2] Y. Hu and A. Ribeiro, "Optimal wireless networks based on local channel state information," *IEEE Trans. Signal Process.*, vol. 60, pp. 4913–4929, September 2012.

[3] A. Ribeiro, "Optimal resource allocation in wireless communication and networking," *EURASIP J. Wireless Commun., Networking*, vol. 2012, August 2012.

[4] Y. Hu and A. Ribeiro, "Adaptive distributed algorithms for optimal random access channels," *IEEE Trans. Wireless Commun.*, vol. 10, pp. 2703–2715, August 2011.

[5] A. Ribeiro, "Ergodic stochastic optimization algorithms for wireless communication and networking," *IEEE Trans. Signal Process.*, vol. 58, pp. 6369–6386, December 2010.

[6] A. Ribeiro and G. Giannakis, "Separation principles in wireless networking," *IEEE Trans. Inf. Theory*, vol. 56, pp. 4488–4505, September 2010.

[7] N. Gatsis, A. Ribeiro, and G. Giannakis, "A class of convergent algorithms for resource allocation in wireless fading networks," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 1808–1823, May 2010.

The separation principles were introduced in [6] and the ergodic stochastic optimization algorithm in [5]. Both of these topics are covered in the tutorial paper [3], which is the best starting point to understand this project. Paper [7] deals with how to solve optimal resource allocation problems in networks with interference limited physical layers. Papers [1], [2], and [4] are the work of Yichuan Hu on understanding optimal wireless networking when terminals have different beliefs about the global state of the network.

B Communication networks for autonomous robot teams

The confluence of advances in robotics and wireless communications has led to the emergence of autonomous robot teams that cooperate to accomplish tasks assigned by human operators. A typical scenario is search and rescue missions in hazardous situations where a team is deployed to scout points of interest. While a designated lead member of the team moves to a specified location, the remaining robots provide mission support. Critical for task accomplishment is the availability of wireless communications. Communication is required to exchange information between robots as well as to relay information to and from the human operators. Availability of wireless communication infrastructure, however, is unlikely in the harsh environments in which autonomous robot teams are to be deployed. Rather, we want the robots to self organize into a wireless network capable of supporting the necessary information exchanges. The goal of this project is to design cyber-physical controllers that determine (physical) trajectories for the robots while ensuring (cyber) availability of communication resources.

The fundamental roadblock to guaranteeing reliable communication is the unpredictability of wireless propagation. The cluttered environments in which robot teams are to be deployed are characterized by severe fading resulting in volatile communications performance even for short distance communications. Channel strength variations in the order of 10 to 20 dB are typical when robots move a distance comparable to the radio wavelength for reference, the wavelength is 30 cm for communications in the 800MHz band.

This drawback notwithstanding, it has to be noted that the goal of the self-organized network is to maintain reliable end-to-end communication between, say, the lead member of the scouting

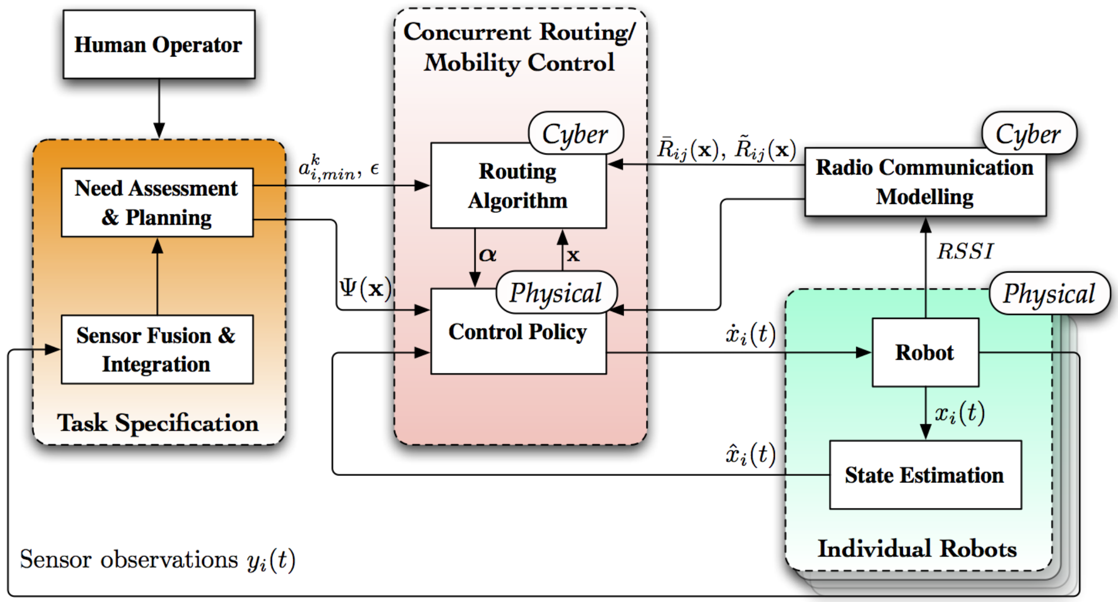

Figure 6: System architecture for joint control of mobility and routing. Task Specification here represents a generic spatial application defined by a convex task potential function $\Psi(x)$ while providing a stream of data to the human operator. Individual robot components consist of the low-level robot control, estimation, and communication. We additionally assume that a subsystem is available to build an online model of radio communication in the environment. The focus of this work is on developing concurrent methods for routing and mobility control.

team and the operation base. We can then exploit spatial redundancy across channels to minimize the effect of point-to-point uncertainty in end-to-end communication rates. By splitting traffic flows between various neighboring robots, we ensure that while failure of a particular link may reduce end-to-end communication rates, it does not interrupt them completely. To realize this idea we adopt a stochastic model of connectivity in which achievable point-to-point rates are random quantities with known mean and variance. If point-to-point rates are random, end-to-end rates are also random. The cyber part of our control algorithms assume given positions and determine routing variables that minimize the probability of end-to-end rates falling below a minimum level of service. The physical control block determines robot trajectories that are restricted to configurations that ensure these probabilities stay above a minimum reliability level.

An architecture diagram is shown in Fig. 6. As with any mobility control system, there is a block performing task specification, a second block executing the control law, and a third block conducting actuation and state estimation. The task specification block interfaces with the human operators and integrates robot observations and requirements to determine specifications that it passes on to the control block. These specifications come in the form of a potential function $\Psi(x)$ that must be minimized and communication rates $a^k_{i,\min}$ that must be maintained at all times. The control block then determines control inputs $\dot{x}(t)$ and network variables $\alpha(t)$ that are conducive to task completion. Individual robots implement the control law $\dot{x}(t)$ and route packets according to variables $\alpha(t)$. Robots also take observations $y_i(t)$, e.g., a video feed, that they relay to task planning and perform position estimation $\hat{x}_i(t)$ that they feedback to the control block. Using available technologies for mapping, control, and state estimation, each robot estimates its position $\hat{x}_i(t)$ and controls its velocity $\dot{x}_i(t)$ with respect to a common known map of the environment. The control algorithm also necessitates access to achievable rates $R_{ij}(x)$ which are provided by a radio communication modeling block. Our work is concerned, for the most part, with the joint selection of suitable communication variables $\alpha(t)$ and path plans $x(t)$. Our work is concerned, for the most part, with the joint control of mobility and routing

In the context of this project we have developed algorithms that rely on the computation of communciation routes and trajectories by a central command center. The plans are then communicated to the individual robots that proceed to implement the plan. A proof of concept of the overall system is shown in this video. Our current effort focuses on the development of distributed algorithms and further experimental validation.

B.1 Decentralized communications protocols

To avoid the computation of routes at a centralized location we are developing algorithms to determine optimal robust routes in a distributed manner. Robots are supposed to have access to local channel information and the ability to communicate with nearby robots but are unaware of the network's topology beyond their local neighborhood. Using local information and variable exchanges between neighbors we devise iterative protocols that determine optimal routing policies as times grows.

A common approach to devise distributed optimal routing algorithms is to implement gradient descent in the dual function of the corresponding optimization problem or subgradient descent if the dual function is nondifferentiable. Optimal routing problems include a maximization objective that determines the metric used to compare different configurations as well as constraints that determine feasible routing variables. To formulate the dual problem we introduce Lagrange multipliers associated with each of the constraints and proceed to find configurations that optimize the linear combination of objective and constraints that they define. The dual problem corresponds to the determination of multipliers that render the constrained optimization and the linear combination optimization equivalent.

The important property of the dual function that makes it appealing to distributed implementations is that it is possible to compute gradients of the dual function in a distributed manner. Robots maintain local routing variables and local multipliers used to enforce their local routing constraints. They then proceed to recursively update primal variables by finding optimal routes for given duals and dual variables by performing a gradient descent step. The first fundamental observation is that dual gradients can be found as the constraint slack associated with the optimal routing configuration. The second fundamental observation is that all of the required computations can be implemented in a distributed manner. Determination of optimal routes requires access to local and neighboring dual variables whereas dual gradients are given as functions of local and neighboring multipliers.

B.2 Decentralized motion planning

For fully autonomous operation, decentralized networking protocols need to be integrated into decentralized motion planning algorithms. Our networking protocols determine routing variables that maximize the probability of survival of end-to-end communication flows. However, for some spatial configurations it is impossible to maintain these probabilities above a minimum target level. Since the likelihood of network survival for those configurations is not sufficient for mission requirements, these spatial arrangements generate a virtual obstacle in configuration space. The purpose of this task is to design decentralized reactive as well as deliberative planners to steer system configurations away from the virtual obstacle region created by network survivability requirements.

In the context of reactive planners the solution involves the incorporation of barrier potentials to avoid crossing into the infeasible space. The key research issue to be addressed here is the integration of the decentralized algorithm for optimal communication with the decentralized version of motion planning. For the development of distributed deliberative planning algorithms the challenge is to let robots make consistent local plans based on perceptions of their different local environments. Robots may know the position of some nearby peers and have good estimates of the communication channels that link them. For some more distant agents this knowledge becomes more uncertain and may even become unavailable for some robots. The foremost difficulty is not the uncertain knowledge but the fact that network state information is different at different robots. Our objective is to determine optimal trajectories that take into account the fact that different terminals have different beliefs on the network state and are bound to select conflicting actions.

The resolution of conflicting actions leads naturally to the use of controllers based on partially observable Markov Decision Processes (POMDP). POMDPs provide a general decision making framework for acting optimally in partially observable domains. In lieu of deterministic state estimates the system's state is described in information space where probability distributions on robots' positions and channels are maintained. Planning decisions are then made to maximize a discounted expected payoff across all possible paths. In a distributed context each robot plans a joint optimal trajectory based on its local belief and proceeds to move according to this plan. Since different robots have different beliefs the actual joint trajectory is different from local beliefs. Upon observation of these different trajectories robots update their local beliefs and repeat the process. The advantage of a POMDP formulation is that it is easy to incorporate probabilistic safety constraints to reduce the likelihood of finding the team in spatial configurations for which it is impossible to guarantee network connectivity.

B.3 Acknowledgments and references

This project got started on collaboration with Dr. Jon Fink, who is now at the Army Research Lab. The work on decentralized algorithms got started on a collaboration with Prof. Mihalis Zavlanos, who is now at Duke University. We also collaborated with Jon's and Mihalis's respective advisors, Prof. Vijay Kumar and Prof. George Pappas. Currently, this project is the Ph.D. work of James Stephan. The following is a list of papers produced in the context of this project:

[1] J. Fink, A. Ribeiro, and V. Kumar, "Algorithms for controlling mobility while maintaining robust wireless connectivity," *IEEE Access*, vol. (submitted), January 2013.

[2] M. Zavlanos, A. Ribeiro, and G. Pappas, "Network integrity in mobile robotic networks," *IEEE Trans. Autom. Control*, vol. 58, pp. 3–18, January 2013.

[3] J. Fink, A. Ribeiro, and V. Kumar, "Robust control for mobility and wireless communication in cyber-physical systems with application to robot teams," *Proc. of the IEEE*, vol. 100, pp. 164–178, January 2012.

[4] J. Fink, A. Ribeiro, and V. Kumar, "Motion planning for robust wireless networking," in *Proc. Int. Conf. Robotics Autom.*, vol. 2419-2426, Saint Paul, MN, May 14-18 2012.

[5] M. Zavlanos, A. Ribeiro, and G. Pappas, "A framework for integrating mobility and routing in mobile communication networks," in *Proc. Asilomar Conf. on Signals Systems Computers*, pp. 1461–1465, Pacific Grove CA, November 6-9 2011.

[6] M. Zavlanos, A. Ribeiro, and G. Pappas, "Distributed control of mobility and routing in networks of robots," in *Proc. IEEE Workshop on Signal Process. Advances in Wireless Commun.*, pp. 236–240, San Francisco CA, June 26-29 2011.

[7] J. Fink, A. Ribeiro, V. Kumar, and B. M. Sadler, "Optimal robust multihop routing for wireless networks of mobile micro autonomous systems," in *Proc. Military Commun. Conf.*, pp. 1268–1273, San Jose CA, October 31 - November 3 2010.

[8] M. Zavlanos, A. Ribeiro, and G. Pappas, "Mobility and routing control in networks of robots," in *Proc. Conf. on Decision Control*, vol. (to appear), pp. 7545–7550, Atlanta GA, December 15-17 2010.

The framework described here is presented in the journal publications [1] and [3]. The work in [3] is more tutorial in nature. Many details are omitted and those are discussed in [1]. A comprehensive presentation of our work on distributed implementations is available in [1]. The conference papers [4] and [7] are preliminary versions of [1] and [3]. The conference papers [5], [6], and [8] are preliminary versions of [2].

C Bayesian network games

In many situations agents in a network want to take actions that are optimal with respect to an unknown state of the world and the actions taken by other agents in the network who themselves are trying to select actions that are optimal in the same sense. In, e.g., trade decisions in a stock market, the payoff that a player receives depends not only on the fundamental (unknown) price of the stock but on the buy decisions of other market participants. After all, the price of a stock moves up as long as it is in high demand, which may or may not be because of sound fundamentals. In such situations players must respond to both, their belief on the price of the stock and their belief on the actions of other players. Similar games can also be used to model the coordination of members of an autonomous team whereby agents want to select an action that is jointly optimal but only have partial knowledge about what the action of other members of the team will be. Consequently, agents select actions that they deem optimal given what they know about the task they want to accomplish and the actions they expect other agents to take.

In both of the examples in the previous paragraph we have a network of autonomous agents intent on selecting actions that maximize local utilities that depend on an unknown state of the world – information externalities – and the also unknown actions of all other agents – payoff externalities. In a Bayesian setting – or a rational setting, to use the nomenclature common in the economics literature – nodes form a belief on the actions of their peers and select an action that maximizes the expected payoff with respect to those beliefs. In turn, forming these beliefs requires that each network element make a model of how other members will respond to their local beliefs. The natural assumption is that they exhibit the same behavior, namely that they are also maximizing their expected payoffs with respect to a model of other nodes' responses. But that means the first network element needs a model of other agents' models which

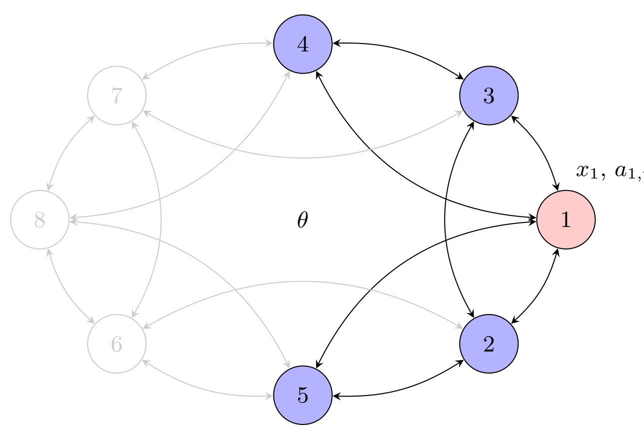

Figure 7: Bayesian network game. Agents want to select actions that are optimal with respect to an unknown state of the world and the actions taken by other agents. We consider repeated versions of this game in which agents observe the actions taken by neighboring agents at a given time.

shall include their models of his model of their model and so on. The fixed point of this iterative chain of reasoning is a Bayesian Nash Equilibrium.

In this project we consider repeated versions of this game in which agents observe the actions taken by neighboring agents at a given time. In observing neighboring actions agents have the opportunity to learn about the private information that neighbors are, perhaps unwillingly, revealing. Acquiring this information alters agents' beliefs leading to the selection of new actions which become known at the next play prompting further reevaluation of beliefs and corresponding actions. In this context we talk of Bayesian learning because the agents' goal can be reinterpreted as the eventual learning of peers' actions so that expected payoffs coincide with actual payoffs. A schematic representation of this model is shown in Fig. 7. The payoff at time t of, say, agent 1 depends on the state of the world θ and the actions $a_{i,t}$ of all other agents. At the initial stage of the game, agent 1 has access to a private signal x_1 which he uses to select the action $a_{1,1}$. It then takes this action and simultaneously observes the actions $a_{i,1}$ of agents in its neighborhood, i.e., $i = 1, 3, 4, 5$. These observed actions reveal further information about θ and the actions that are to be taken by other agents at time $t = 2$. This prompts selection of a likely different action $a_{1,2}$. Observation of actions $a_{i,2}$ of neighboring agents reveals further information on the world state and actions of other players and prompts further reevaluation of the action to be taken at the subsequent stage.

This project involves two relatively separate thrusts respectively concerned with the asymptotic properties of the Bayesian Nash Equilibria and with the development of algorithms that agents can use to compute their equilibrium actions. The former thrust is conceptual and mostly of interest as a model of human behavior in social and economic networks. The latter thrust is of interest to enable the use of network games as distributed algorithms to plan the actions of members of an autonomous team.

C.1 Asymptotic learning

As mentioned earlier we can reinterpret an agent's purpose as the eventual learning of all available information about the state of the world and the peers' actions. The question we try to answer in this thrust is whether this information is eventually learnt or not.

Different behavioral assumptions lead to different outcomes. In particular, the way agents revise their views in face of new information and the actions they choose given these views determines the long run outcome of the game. In our preliminary work we assume that agents are myopic in that they choose actions at each stage of the game which maximize their stage payoffs, without regard for the effect of these actions on their future payoffs. We use this behavioral assumption to prove formal results regarding the agents asymptotic equilibrium behavior. Our analysis yields several interesting results. First, agents actions asymptotically converge for almost all realizations of the game. Furthermore, given a connected observation network, agents actions converge to the same value. In other words, agents eventually coordinate on the same action. Second, agents reach consensus in their best estimates of the underlying parameter. Finally, if some agent can eventually observe her own realized payoffs, agents coordinate on an action which is optimal given the information dispersed among them. This result suggests that in a coordination game – where agents interests are aligned – repeated interactions between agents who are selfish and myopic could eventually lead them to coordinate on the socially optimal outcome.

We are currently investigating further properties of steady state equilibrium plays for Bayesian games with myopic players. We are also looking into networks with non-myopic players and relaxing assumptions on the knowledge of network topology and other structural information.

C.2 Quadratic network game filter

The burden of computing a Bayesian Nash equilirium in repeated games is, in general, overwhelming even for small sized networks. This intractability has led to the study of simplified models in which agents are non-Bayesian and update their beliefs according to some heuristic rule. A different simplification is obtained in models with pure information externalities where payoffs depend on the self action and an underlying state but not on the actions of others. This is reminiscent of distributed estimation since agents deduce the state of the world by observing neighboring actions without strategic considerations on the actions of peers. Computations are still intractable in the case of pure information externalities and for the most part only asymptotic analyses of learning dynamics with rational agents are possible. Explicit methods to maximize expected payoffs given all past observations of neighboring actions are available only when signals are Gaussian or when the network structure is a tree. For the network games considered in this project in which there are information as well as payoff externalities, not much is known besides the asymptotic analyses of learning dynamics described above.This is an important drawback for the application of network games to the implementation of distributed actions in autonomous teams. The purpose of this thrust is to develop algorithms to enable computation of equilibrium actions.

Our first result considers payoffs represented by a utility function that is quadratic in the actions of all agents and an unknown parameter. At the start of the game each agent makes a private observation of the unknown parameter corrupted by additive Gaussian noise. To determine a mechanism to calculate equilibrium actions we introduce an outside clairvoyant observer that knows all private observations. For this clairvoyant observer the trajectory of the game is completely determined but individual agents operate by forming a belief on the private signals of other agents. We start from the assumption that this probability distribution is normal with an expectation that, from the perspective of the outside observer, can be written as a linear combination of the actual private signals. If such is the case we can prove that there exists a set of linear equations that can be solved to obtain actions that are linear combinations of estimates of private signals. This result is then used to show that after observing the actions of their respective adjacent peers the probability distributions on private signals of all agents remain Gaussian with expectations that are still linear combinations of the actual private signals. We proceed to close a complete induction loop to derive a recursive expression that the outside clairvoyant observer can use to compute BNE actions for all game stages. We then leverage this recursion to derive the Quadratic Network Game (QNG) filter that agents can run locally, i.e., without access to all private signals, to compute their equilibrium actions. A schematic representation of the QNG filter is shown in Fig. 8 to emphasize the parallelism with the Kalman filter. The difference is in the computation of the filter coefficients which require the solution of a system of linear equations that incorporates the belief on the actions to be taken by other agents.

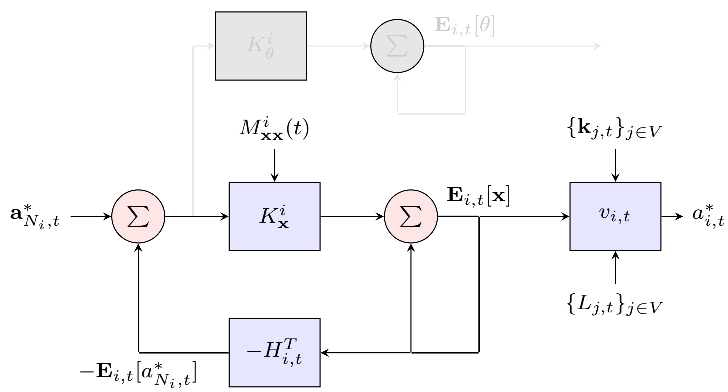

Figure 8: Quadratic Network Game (QNG) filter. Agents tun the QNG filter to compute Bayesian Nash equilibrium actions in games with quadrate payoffs and Gaussian private signals.

Our current research program includes developing analogues of the QNG filter for different network types and leveraging the QNG filter to solve games where payoffs are not quadratic and private signals are not Gaussian. The similarity with Kalman filters is instrumental here as it points out to the existence of filters akin to extended or unscented Kalman filters as well as particle filters.

C.3 Acknowledgments and references

This is the Ph.D. work of Ceyhun Eksin and Pooya Molavi. We also count on invaluable help from Ali Jadbabie. The list of publications resulting from this project is the following:

[1] C. Eksin, P. Molavi, A. Ribeiro, and A. Jadbabaie, "Learning in network games with incomplete information," *IEEE Signal Process. Mag.*, vol. (to appear), May 2013.

[2] C. Eksin, P. Molavi, A. Ribeiro, and A. Jadbabaie, "Bayesian quadratic network game filters," *IEEE Trans. Signal Process.*, vol. (submitted), January 2013.

[3] C. Eksin, P. Molavi, A. Ribeiro, and A. Jadbabaie, "Bayesian quadratic network game filters," in *Proc. Int. Conf. Acoustics Speech Signal Process.*, vol. (submitted), Vancouver Canada, May 26-31 2013.

[4] C. Eksin, P. Molavi, A. Ribeiro, and A. Jadbabaie, "Dynamic games with side information in economic networks," in *Proc. Asilomar Conf. on Signals Systems Computers*, vol. (to appear), Pacific Grove CA, November 4-7 2012.

[5] C. Eksin, P. Molavi, A. Ribeiro, and A. Jadbabaie, "Learning in linear games over networks," in *Proc. Allerton Conf. on Commun. Control Computing*, vol. (to appear), Monticello IL, October 1-5 2012.

To understand this project the best place to start is the tutorial article [1] which appeared in the Signal Processing Magazine. For complete details on the QNG filter the journal submission [2] os the best place to start. The conference submission in [3] is an abridged version of the journal paper [2]. Conference submissions [4] and [5] contain our results on asymptotic learning.

D Circles of trust: Hierarchical clustering in asymmetric networks

Miranda trusts Billy who trusts Ariel who trusts Miranda, but there has not been enough interactions in the opposite direction to establish trust. When these three people meet, shall they trust each other? I.e., are they part of a circle of trust? The objective of this project is to develop an axiomatic theory to provide an answer to this question. In general, we start with a network were nodes represent individuals and *directed* edges represent a trust dissimilarity from the originating node to the end node. Small values of this dissimilarity signify large amounts of trust of the edge's source node on the edge's destination. Our goal is the study of the formation of trust groups in the network. I.e., the determination of the level of trust at which two individuals are integrated in a trust cluster given not only their direct interactions but their indirect interactions through other members of the network. It may make sense for Miranda, Billy, and Ariel to trust each other, because they all either trust each other directly, or have trust on someone that trusts the person they don't know.

Once the problem is written in this language it is clear that determining circles of trust is akin to finding clusters in an asymmetric network for a given resolution level. The determination of a family of clusterings indexed by this resolution parameter is a problem known as hierarchical clustering. Simple as this sounds, the problem is that clustering in general and clustering using asymmetric data in particular is a poorly understood problem. There are plethora of methods that can be chosen to perform clustering, but these methods are based on heuristic intuition, not fundamental principles. Beyond purist concerns, lack of theoretical understanding is also a practical problem for clustering of asymmetric data because the intuition backing clustering methods is drawn from geometric point clouds. This intuition does not carry when the given data is not metric as in the case of asymmetric trust dissimilarities. E.g., in the network in Fig. 18 nodes a and b are closest to each other clockwise, but farthest apart counterclockwise, c and d seem to be closest on average, yet, it seems that all nodes are relatively close as it is possible to circle the graph clockwise without encountering a dissimilarity larger than 3.

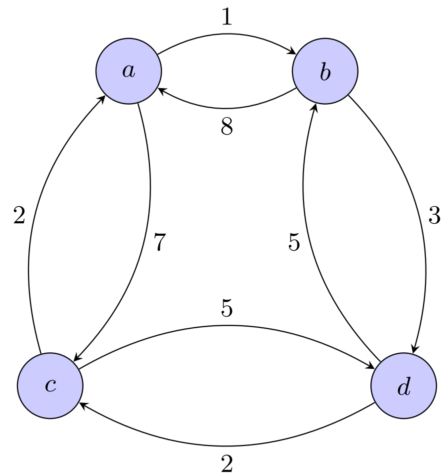

Figure 9: Trust network. Arcs denote trust dissimilarity, e.g., a has substantial trusts in b, but b has little trust in a. Clustering intuition is precarious. Nodes seem close clockwise but far apart counterclockwise.

Even though asymmetric clustering intuition is difficult in general, there are some particular specks of intuition that we can exploit to gain insight into the general problem. These intuitive statements can be postulated as axioms that restrict the space of allowable asymmetric hierarchical clustering methods. To the extent that the axioms are true, the properties of this reduced space of methods are fundamental properties of asymmetric hierarchical clustering and by extension fundamental properties of the formation of circles of trust. This is the approach advocated in this project.

D.1 Axioms of value and transformation

In our preliminary investigations we have postulated three axioms that we call the axioms of value, influence, and transformation. These axioms are stated formally in our published work but they correspond to the following intuitions:

(A1) Axiom of Value. For a network with two nodes the nodes are clustered together at the resolution at which both trust each other, namely, the maximum of the two trust dissimilarities between them.

(A2) Axiom of Transformation. If we consider a network and reduce all pairwise trust dissimilarities, the level at which two nodes become part of the same circle of trust is not larger than the level at which they were clustered together in the original network.

Axiom (A1) says that in a network with two nodes p and q and dissimilarities $A_X(p,q) = \alpha$ and $A_X(q,p) = \beta$, the nodes are reported as separate clusters for resolutions $\delta < \max(\alpha, \beta)$ and as a single cluster for resolutions $\delta \geq \max(\alpha, \beta)$. This is reasonable because at resolutions $\delta < \max(\alpha, \beta)$ one node can influence the other but not vice versa, which in most situations means that the nodes are not alike since trust flows in a single direction. Axiom (A2) states that increasing the level of trust

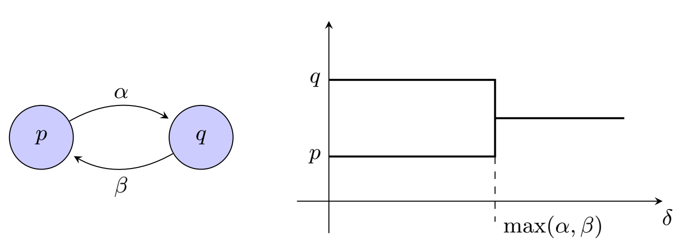

Figure 10: Axiom of Value. For a two node network nodes are clustered together at the resolution at which both can influence each other.

between some nodes may result in the formation of additional circles of trust but cannot result in the dissolution of existing circles. This is erasable because a reduction in trust dissimilarities is expected to generate more trust in the network.

Despite their apparent weakness, axioms (A1)-(A2) are a source of strong structure. Our first preliminary result is the derivation of two asymmetric hierarchical clustering methods that abide to these axioms. The first method insists that trust propagate only through arcs in which there is bidirectional trust and is therefore termed reciprocal clustering. The second method allows trust to propagate unidirectionally and is thus termed nonreciprocal clustering. That these methods comply with (A1)-(A2) is not particularly surprising. However, we have proved that any clustering method that satisfies axioms (A1)-(A2) lies between reciprocal and nonreciprocal clustering in a well defined sense. Specifically, any clustering method that satisfies axioms (A1)-(A2) forms circles of trust at resolutions larger than the resolutions at which they are formed with nonreciprocal clustering, and smaller than the resolutions at which they are formed with reciprocal clustering. These preliminary result endows reciprocal and nonreciprocal clustering with special meaning. For a given resolution level, nodes that *do not* cluster together with nonreciprocal clustering cannot be part of a circle of trust. Nodes that *do* cluster together with reciprocal clustering are definitely part of a circle of trust. In between, the answer depends on the extent to which reciprocal trust propagation is required or nonreciprocal trust propagation is acceptable.

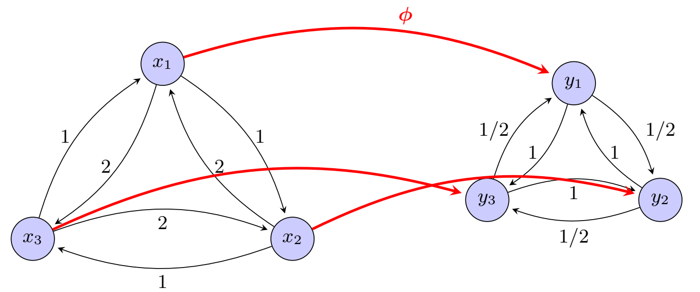

Figure 11: Axiom of Transformation. A dissimilarity reducing map $\phi : X \to Y$ produces dendrograms where clusters in the original network may be combined in the transformed network but cannot be separated.

D.2 Ongoing work

These preliminary results are enticing but far from a complete axiomatic exploration on the formation of circles of trust. To enrich this exploration we are currently pursuing four interrelated research thrusts that we preview in the following.

Circles of trust. Nonreciprocal and reciprocal clustering set lower and upper bounds in the formation of circles of trust providing the basis for their empirical study in social networks. This thrust encompasses a data collection effort and corresponding data analysis. Questions of interest include the symmetry and segmentation of trust. We associate symmetry of trust with the difference between reciprocal and nonreciprocal clusters. We identify segmentation of trust as the number of separate clusters as a function of the resolution level.

Further axiomatic constructions. Reasonable though they are, axioms (A1)-(A2) are a particular selection. This research thrust leverages the developed techniques to study additional axiomatic constructions. We consider different versions of axioms (A1) and (A2) that could lead to more general, more restricted, or simply different sets of admissible clustering methods. We are also studying additional axioms that can be added on top of (A1)-(A2) with the objective of finding a set of axioms leading to uniqueness results.

Stability. Two networks that are close to each other shall yield hierarchical clusters that are also close to each other. While this sense of stability looks like, and should be, a precondition for practical applicability, many (symmetric) clustering algorithms used in practice are not stable in this sense. Studying stability in asymmetric clustering is further complicated by the nonexistence of suitable tools to formalize the proto-continuity statement of "networks close to each other." The first objective of this research thrust is to define generalized versions of Gromov-Hausdorff and Gromov-Wasserstein distances that can be used to formalize a notion of continuity in the space of networks. With this tool in hand we derive conditions to guarantee that networks arbitrarily close to each other yield hierarchical clusters that are also arbitrarily close to each other.

Intermediate clustering methods. Nonreciprocal and reciprocal clustering lower and upper bound the range of hierarchical clustering methods. A first question that arises is if there really are clustering methods lying between these two. We believe the answer to that question is probably positive. The reason for this belief is that the outputs of hierarchical clustering algorithms are trees. If we have two trees that satisfy axioms (A1)-(A2) we can combine them with each other by pruning branches that we then graft in the other tree. It seems reasonable that we could accommodate this grafting process so that axioms (A1)-(A2) are satisfied in the rearranged tree. We are currently studying this and other techniques that may lead to intermediate clustering methods.

D.3 Acknowledgments and references

This is the Ph.D. work of Santiago Segarra. This project builds on the work of Prof. Facundo Memoli and Prof. Gunnar Carlsson on hierarchical clustering for metric data. We are glad to count them as collaborators in this project which wouldn't be feasible without their help. Some preliminary results have appeared in the following paper:

[1] G. Carlsson, F. Memoli, A. Ribeiro, and S. Segarra, "Axiomatic construction of hierarchical clustering in asymmetric networks," in *Proc. Int. Conf. Acoustics Speech Signal Process.*, (to appear), Vancouver Canada, May 26-31 2013.

E Distributed optimization

In distributed optimization problems agents want to find local variables that are optimal with respect to a local utility while satisfying linear coupling constraints with variables of neighboring nodes. For a more precise definition consider a network composed of N nodes indexed by i and denote as $n(i)$ the set of neighbors of i and as $\mathbf{C}$ the directed edge incidence matrix of the graph. Each of the nodes keeps track of a local variable x_i and a local concave utility function $f_i(x_i)$. We want to select variables that maximize the sum utility $\sum_i f_i(x_i)$ subject to the requirement that the variables x_i are such that $x_i = x_j$ for all nodes $j \in n(i)$ in its neighborhood. For a connected network this is the same as requiring that the variables x_i have the same value for all nodes and this particular problem formulation is therefore termed a consensus optimization problem. Defining the vector $\mathbf{x} = [x_1, \ldots, x_N]^T$ grouping all local variables and letting $\mathbf{0}$ be the all-zero vector of proper dimension we can write the consensus optimization problem as

$$\mathbf{x}^* = \arg\max \sum_i f_i(x_i), \quad \text{s.t. } \mathbf{C}^T\mathbf{x} = \mathbf{0}. \tag{2}$$

A particular application of a consensus optimization problem is distributed maximum likelihood estimation. In such case the variables x_i take the place of a local estimate, the functions $f_i(x_i)$ represent local log likelihoods determined by local observations, and the goal is for each agent to compute a local estimate using global information. Another important application can be obtained by replacing the matrix $\mathbf{C}$ in (2) to model optimal flow problems.

Regardless of the particular application it is of interest to derive distributed optimization algorithms to solve distributed optimization problems. These algorithms rely on iterative computation of candidate solutions $x_i(t)$ and variable exchanges between neighboring nodes so that as time progresses the iterates $x_i(t)$ approach the optimal variable x_i^*. The original use of distributed optimization algorithms was for distributed control and information aggregation in wireless sensor networks. Currently, there is renewed interest in these algorithms as per their applicability to solve massive dimensional optimization problems in server clusters. In the latter setting there is interest in subdividing the optimization problem into separate servers while keeping communication requirements subdued.

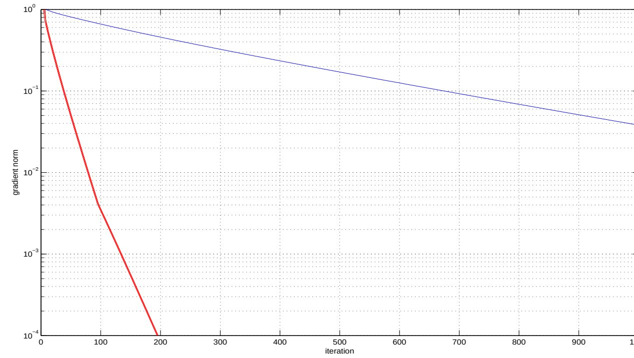

Figure 12: Infinity norm distance from primal iterates to optimal arguments for ADD-1 (red) and ADMM (blue). The curvature correction of ADD retains reasonable convergence times for problems that are not well conditioned. (Directed cycle network with 50 nodes and 50 random edges. Quadratic primal objectives with condition number 10.)

There are different approaches that lead to distributed optimization algorithms. Our work concentrates on methods that operate in the dual domain. Without digressing into technical details the structure of the dual function is such that it is possible to compute gradients of the dual function in a distributed manner. These gradients can then be used to implement a dual gradient descent algorithm that converges toward the optimal dual variables from which the optimal primal variables of the optimization problem introduced above can be recovered as a byproduct. Dual gradient descent algorithms are valuable for their simplicity but their application is limited to problems whose dual functions are well conditioned. Since the condition number of the dual function is roughly given by the product of the condition number of the primal function and the ratio between the largest and smallest eigenvalues of the graph's Laplacian, this requires tame networks and primal functions. Tame primal functions are those with good condition numbers. Tame networks are those that have small diameter and some form of regularity on the number of connections per node.

A similar algorithm with better convergence properties is the Alternating Direction Method of Multipliers (ADMM) which is based on the addition of a quadratic regularization term to the the problem's Lagrangian. The ADMM is less sensitive to the condition of the dual function. It's performance still degrades with increasing condition number but the degradation is much smaller than the degradation corresponding to dual gradient descent. Whereas dual descent is all but impractical for problems having more than a few nodes and condition numbers not close to one, ADMM is slow but acceptable for problems with moderate number of nodes and moderate condition numbers. A convergence rate illustration is shown in Fig. 12 that depicts the infinity norm distance between ADMM iterates and the primal optimal argument $\mathbf{x}^*$. The network is a directed cycle graph with 50 nodes to which 50 random edges have been added. The primal objective is a quadratic function with condition number 10. A distance to optimality of about 4×10^{-2} is attained in 10^3 iterations.

While better behaved than dual descent, convergence of ADMM is still inadequate and can be made arbitrarily slow by increasing the condition number of the primal function or modifying the network to increase the ratio between the largest and smallest eigenvalue of the graph's Laplacian. Ultimately, this drawback can only be corrected by implementing curvature corrections as in Newton and quasi-Newton methods. Our research on this project is concerned with the development of distributed optimization methodologies that emulate Newton's method as a means of attaining quadratic convergence rates. We discuss these methods in the following section but a preliminary simulation is shown in shown in Fig. 12 depicting the infinity norm distance between iterates obtained by our distributed optimization methods and the primal optimal argument $\mathbf{x}^*$. For the same problem parameters used for ADMM, we converge to distance 10^{-2} in less than 100 iterations and to distance 10^{-4} in 200 iterations.

E.1 Accelerated dual descent

The natural alternative to accelerate convergence rate of first order gradient descent methods is to use second order Newton methods, but these methods cannot be implemented in a distributed manner. Indeed, since (centralized) implementation of the Newton method necessitates computation of the inverse of the dual Hessian, it follows that a distributed implementation would require each node to have access to a corresponding row of this inverse Hessian. It is not difficult to see that the dual Hessian is in fact a weighted version of the networks Laplacian and that as a consequence its rows could be locally computed through information exchanges with neighboring nodes. Its inversion, however, requires global information.

The insight at the core of this research thrust is to consider a Taylor's expansion of the inverse Hessian which, being a polynomial with the Hessian matrix as variable, can be implemented through local information exchanges. More precisely, considering only the zeroth order term in the Taylor's expansion yields an approximation to the Hessian inverse based on local information only. The first order approximation necessitates information available at neighboring nodes and in general, the Nth order approximation necessitates information from nodes located N hops away. The resultant family of Accelerated Dual Descent (ADD) algorithms permits a tradeoff between accurate Hessian approximation and communication cost. We use ADD-N to represent the N th member of the ADD family. ADD-N uses the Nth order Taylor approximation of the Hessian inverse by collecting information from terminals N hops away.

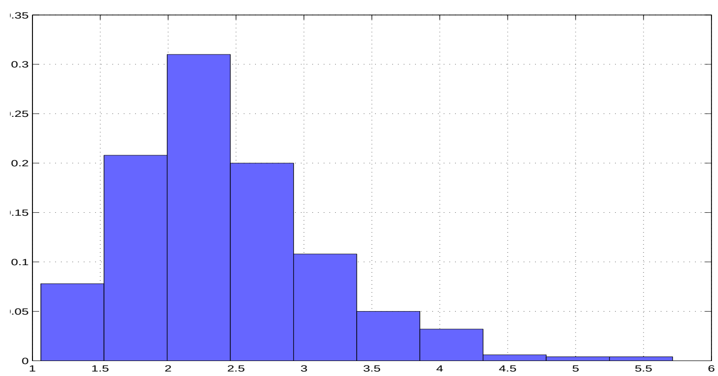

Figure 13: Relative convergence times of ADMM with respect to ADD-1 for a well conditioned problem. Convergence times are similar since ADD's curvature correction is of little help in a problem that already has a tractable shape. (Directed cycle network with 20 nodes and 40 random edges. Quadratic primal objectives with condition number 1. Time to reach infinity norm distance 10^{-4} to optimal arguments)

Our research on ADD has shown that convergence of algorithms in the ADD family follows three distinct phases. The behavior during the first phase is defined by an approximate backtracking line search which ensures a constant decrease towards the optimal objective. During the second phase ADD-N algorithms exhibit quadratic convergence despite the fact that they rely on approximate Newton directions. In the third and terminal phase convergence slows down to linear as errors in the Newton step become comparable to the steps magnitude. The first two phases are akin to the linear and quadratic phases of Newton algorithms. The transition between the second and third phase can be delayed by increasing N, although this may result in increased overall communication cost. Our numerical studies corroborate substantial imporvements in convergence times. Fig. 13

shows the ratio of times required by ADMM and ADD-1 to reach an infinity norm distance to the optimal arguments of at least 10^{-4} for a primal quadratic objective with condition number 1 in a directed cycle network with 20 nodes to which 40 random edges have been added. Since the curvature of the original problem is benign both methods exhibit similar convergence rates. Modifying the objective to yield a condition number of 10 results in convergence times for ADD-1 that are 20 times faster than the convergence times of ADMM as we show in Fig. 14.

Our current work includes extensions of ADD to problems with dual functions that have singular Hessians and that are not differentiable at all times. We are also considering to stochastic optimization problems and alternative approaches based on different matrix splittings and quasi-Newton methods. We also have a parallel research project in which we consider dynamic optimization problems characterized by time varying objectives. In the latter case the challenge is that in the time it takes for iterative algorithms to converge to the optimal argument the objective has changed enough for the optimal operating point to be different.

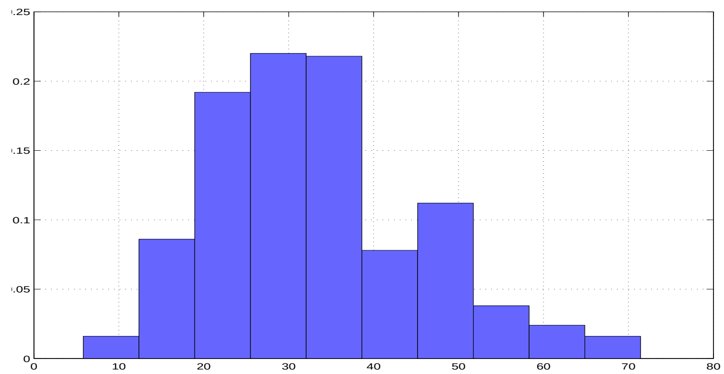

Figure 14: Relative convergence times of ADMM with respect to ADD-1 for an ill conditioned problem. Convergence times are more than one order of magnitude faster for ADD. The curvature correction is essential to achieve reasonable convergence times. (Quadratic primal objectives with condition number 1. Other parameters as in Fig. 13.)

E.2 Acknowledgments and references

My work on distributed optimization got started on a collaboration with Ioannis Schizas and Georgios Giannakis where we adapted the ADMM to solve distributed maximum likelihood estimation problems. As an aside note, observe that while interest in the ADMM for distributed network optimization has been revived by an excellent tutorial manuscript by Boyd et al., it is Ioannis Schizas who deserves credit for adopting the ADMM as a tool for solving distributed consensus optimization problems. This paper predates Boyd et al. by a good three years and spawned a literature of a couple hundred papers dealing with various applications.

The Accelerated Dual Descent (ADD) family of algorithms is the Ph.D. work of Mike Zargham in collaboration with Ali Jadbabaie, Asuman Ozdaglar, and myself. The list of papers related to this work is the following:

[1] M. Zargham, A. Ribeiro, A. Jadbabaie, and A. Ozdaglar, "Accelerated dual descent for network optimization," *IEEE Trans. Autom. Control*, vol. (revised), August 2012.

[2] M. Zargham, A. Ribeiro, and A. Jadbabaie, "Network optimization under uncertainty," in *Proc. Conf. on Decision Control*, vol. (to appear), Maui Hawaii, December 10-13 2012.

[3] M. Zargham, A. Ribeiro, and A. Jadbabaie, "A distributed line search for network optimization," in *Proc. American Control Conf.*, pp. 472–477, Montreal Canada, June 27-29 2012.

[4] M. Zargham, A. Ribeiro, A. Ozdaglar, and A. Jadbabaie, "Accelerated dual descent for network optimization," in *Proc. American Control Conf.*, pp. 2663–2668, San Francisco CA, June 29 - July 1 2011.

The development of the ADD family is presented in the journal paper [1]. Conference papers [3] and [4] contain preliminary abridged versions of the material that went into the preparation of [1]. Conference submission [2] is a preliminary report on the use of ADD for stochastic optimization problems.

The somewhat parallel work on dynamic distributed optimization is part of the Ph.D. work of Felicia Jakubiec. Qing Ling from the University of Science and Technology of China (USTC) has contributed his invaluable expertise during his visit to Penn. Our work on this topic has appeared here:

[1] F. Jakubiec and A. Ribeiro, "D-MAP: Distributed maximum a posteriori probability estimation of dynamic systems," *IEEE Trans. Signal Process.*, vol. 61, pp. 450–466, February 2013.

[2] F. Jakubiec and A. Ribeiro, "Distributed maximum a posteriori probability estimation for tracking of dynamic systems," in *Proc. Asilomar Conf. on Signals Systems Computers*, vol. (to appear), Pacific Grove CA, November 4-7 2012.

[3] F. Jakubiec and A. Ribeiro, "Distributed maximum a posteriori probability estimation of dynamic systems with wireless sensor networks," in *Proc. Int. Conf. Acoustics Speech Signal Process.*, pp. 2857–2860, Kyoto Japan, March 25-30 2012.

Our work on dynamic optimization is covered in the journal submission [1]. Conference papers [2] and [3] are abridged versions of [1].

F Wireless Control Systems

In this project we want to study networked control systems characterized by the separation of sensing and actuation in different physical devices with control loops incorporating the communication of plant state information over a wireless channel. When sensors and controllers communicate over a wireless channel the cost of controlling the plants gets mixed with the cost of sending plant state information from sensors to controllers. The more information the sensors convey the more precise actuation becomes, but the resulting increase in power consumption at the sensors leads to rapid depletion of its energy resources. It is therefore apparent that a tradeoff emerges between plant performance and power consumption. To quantify this tradeoff we study the problem of selecting plant inputs and power management policies that are jointly optimal with respect to a cost that accounts for plant regulation costs and the costs of conveying information from sensors to controllers.

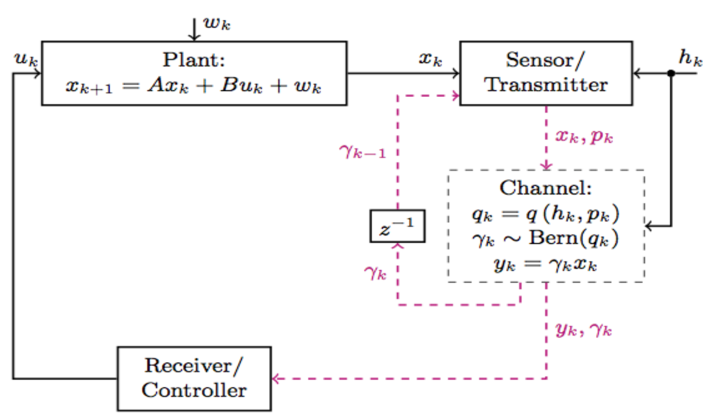

Figure 15: Wireless control system architecture. The sensor measures plant state x_k and fading channel h_k and transmits with power p_k. Messages are decoded at the controller with probability q_k that depends on the channel state h_k and the power p_k. We want to find policies that are jointly optimal with respect the aggregate costs of plant regulation and information transmission from sensors to controllers.

A popular alternative to regulate communication cost in wireless control systems is the notion of event triggered control where the idea is to prolong the time elapsed between successive communications by avoiding transmission as long as the plant performance does not deteriorate much. Such triggering of transmissions implicitly reduces the communication cost but transmission expenses are not explicitly accounted. A related approach is to assign a fixed cost to each transmission attempt and proceed to minimize the combination of a control error cost and an aggregate communication penalty. Optimal transmission policies that minimize the aggregate estimation error and communication cost can be characterized using a formulation in terms of infinite horizon Markov decision processes.

Event triggered control and communication penalty formulations are sufficient in some settings but they do not permit the modeling of fading effects neither allow flexibility in the allocation of power to protect some transmissions more than others. Furthermore, it is not possible to incorporate contention effects into either formulation making it difficult to consider cases in which multiple sensors and multiple controllers operate over a shared wireless channel. In this project we model communication cost as the transmitted power at the sensor side. Powers are selected as a function of the plant state and the fading channel realization and affect the likelihood of successful packet decoding through a known complementary error function. The transmitted power is combined with the conventional linear quadratic regulator (LQR) cost to define an aggregate infinite horizon cost that we seek to minimize through proper joint selection of plant and power control policies. Taking advantage of this problem formulation we can mitigate fading through power adaptation to channel conditions and adapt transmitted power to plant conditions so as to, e.g., increase the likelihood of successful packet decoding when the plant state deviates from target. Channel contention is also easily incorporated by modifying the complementary error function to account for packet losses due to simultaneous transmissions.

F.1 Single plant over point-to-point channel

In the context of this project we have begun by studying the (simple) case in which we have a single control loop closed over a point to point channel. The corresponding system architecture is shown in Fig. 15. A sensor measures the plant state x_k and the fading channel state h_k and transmits with power p_k. Messages are successfully decoded at the controller with a probability q_k that depends on the channel state h_k and the power p_k. If the information is successfully decoded by the controller he learns the current plant state. Otherwise, he propagates past estimates to update his belief on the plant state. In either case a control input u_k is computed and applied to the plant. The goal is to find control inputs u_k and power controls p_k that are good in terms of keeping the state x_k close to target and the power consumption p_k small.

Conceptually, we expect the target successful decoding probability q_k to be close to 1 when the channel and plant state are large and close to zero when they are small. A conceptual colormap of target decoding probability versus channel realization in the horizontal axis and plant state in the vertical axis is shown in Fig. 16-left. When the channel realization is close to zero, successful communication requires investment of a significant amount of power. Thus, the sensor abstains from transmitting unless it is indispensable due to the plant state being far from target – southeast region of the plot. When the channel realization is large transmission is cheap in terms of power and a message is sent even if the plant is close to target – northwest region of the plot. When the channel realization as well as the plant state are large, the target decoding probability is close to 1 because it is both, cheap and necessary – northeast region of the plot. When channel and state are small the target decoding probability is zero because transmission is expensive and unnecessary – sowthwest region of the plot.

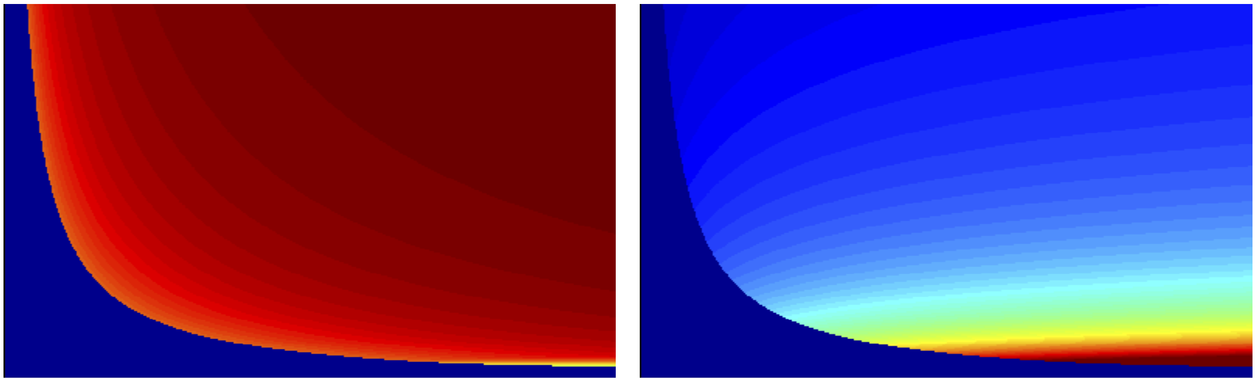

Figure 16: Decoding probability (left) and transmitted power (right) as a function of channel (horizontal axis) and plant (vertical axis) states. When channel and state are small the target decoding probability and the transmitted power are zero because transmission is expensive and unnecessary – sowthwest. When both are large target decoding probability is null and transmit power small – northeast. For large channel and small state transmission doesn't help much but is cheap – northwest. For small channel and large plant transmission is expensive but necessary – southeast.

The corresponding conceptual power map is shown in Fig. 16-right. When channel and state are small the transmitted power are zero because transmission is expensive and unnecessary – sowthwest region go the plot. When both are large a small amount of power is invested to yield a large decoding probability that is instrumental in bringing the plant closer to target – northeast region go the plot. For large channel and small state transmission doesn't help much but is cheap and a small amount of power is invested in a transmission attempt – northwest region go the plot. When the channel realization is small and the plant is far from traget transmission is expensive but necessary resulting in the investment of a significant amount of power in the transmission attempt – southeast region go the plot.

The conceptual allocation in Fig. 16 has been corroborated by our analysis and algorithmic development. We have identified a restricted information structure that permits decoupling of optimal plant and power control policies. For this particular information structure the usual LQR control law becomes optimal at the controller side while the optimal communication policy follows from a Markov decision process (MDP) formulation involving transmitted power and the state estimation error at the controller side. We then leverage this separation principle to express optimal power control policies in terms of a value function. While this does not allow computation of optimal policies it does allow us to understand their qualitative properties. These qualitative properties are the source of the target decoding probability map in Fig. 16 and the transmitted power map in Fig. 16.

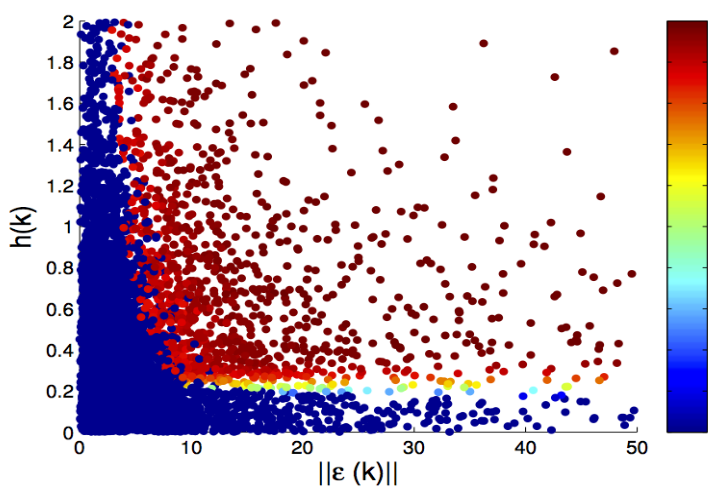

Figure 17: Power allocation of a rollout policy for joint optimization of control inputs and transmitted power. Allocated power is consistent with the conceptual plot in Fig. 16-right.

We can think of the event-triggered paradigm as a special case in our formulation, where instead of deciding how much power to allocate to the transmission attempt we just decide whether to transmit or not. This interpretation is fostered by the realization that conventional event triggered policies emerge as the optimal communication strategy if the sensor uses capacity achieving forward error correcting codes. We have also derived suboptimal power control policies using a rollout algorithm. Numerical simulations of this algorithm yield power allocations like the one shown in Fig. 17, which is consistent with the conceptual plot in Fig. 16-right.

F.2 Ongoing work, acknowledgements, and references

Most of our current effort is centered on developing generalizations where there are multiple sensors and controllers sharing a common wireless channel. We are considering different contention protocols which give rise to different policies to manage channel scheduling along with power allocation and plant control policies. We are also working on alternative formulations where we account not only for the transmitted power at the sensor but also by standby power consumption at the receiver and transmitter ends. In this case sleeping becomes important and leads to policies where sleeping times needs to be prearranged based on beliefs about the evolution of the plant state.

This is the Ph.D. work of Konstantinos Gatsis who should the one to congratulate on whatever is good about the results obtained in the context of this project. We also count on the substantial expertise of Dr. Miroslav Pajic and Prof. George Pappas. The list of publications resulting from this project is the following:

[1] K. Gatsis, A. Ribeiro, and G. Pappas, "Optimal power management in wireless control systems," *IEEE Trans. Autom. Control* (submitted), October 2012.

[2] K. Gatsis, A. Ribeiro, and G. Pappas, "Optimal power management in wireless control systems," in *Proc. American Control Conf.* (to appear), Washington DC, June 17-19 2013.

If you want to understand our work here, the journal paper submission [1] is the place to start. The conference paper [2] is an abridged version of [1] and is a finalist for the best student paper award the 2013 American Control Conference (ACC).

G What's in Shakepeare's name? Authorship attribution from word adjacency networks

For more than a century several crank theories have been popular in some literary circles stating that he whom we call Shakespeare didn't write any of the plays for which he is famous. There are some who believe that Shakespeare's plays were written by Francis Bacon, while some others believe the plays were written by Christopher Marlowe, and yet another group who thinks that Edward de Vere, Earl of Oxford, is the author of the plays that we attribute to William Shakespeare of Stratford upon Avon. This latter theory has recently captivated popular imagination upon the release of the movie "Anonymous." Although these alternative Shakespearean authorship hypotheses have had prominent advocates throughout history reputable Shakespearean scholarship gives them little credence. Nevertheless, there are many open questions about the authorship of some Shakespearan plays going from the uncertain attribution of the "The Taming of the Shrew" to the almost certain fact that some plays where written in collaboration with contemporary authors.

There are, therefore, several important questions regarding Shakespearean plays: (i) Who wrote these plays. (ii) Were these plays written by a single person? (iii) If there are collaborators, who were they? To put it in Shakespearean language the question is: What's in Shakepeare's name?

Shakespeare himself, whomever he, she, or they were, provided an answer to the most fundamental aspect of this question:

> *"What's in a name? That which we call a rose*
> *by any other word would smell as sweet."*

In the end, it doesn't matter if "The Tempest" was written by William Shakespeare, Francis Bacon, Christopher Marlowe, Edward de Vere, or a team of anonymous writers. It would still be a beautiful play and two of my children would still be called Miranda and Ariel. Yet, the third one may not had been named Guille (the Spanish language version of William). His name could had been Francisco (Francis), Cristobal (Christopher), or Eduardo (Edward) instead. This would be catastrophic in many accounts, but it could be worse. It is possible that he should have multiple names some of which may not be known at all. On a more important note, it is impossible to get a researcher to waste enough time on a question he believes he can answer. The purpose of this project is therefore to answer the question: What's in Shakespeare name? Or to put in less poetic terms our goal is to develop methods for authorship attribution. Besides its use in answering authorship questions about playwrights of the English Renaissance, the tools we are developing are applicable in data forensics as well as well as to the detection of plagiarism and other forms of academic malfeasance.

The goal of authorship attribution is to match a text of unknown or disputed authorship to one of a group of potential candidates. More generally, it can be seen as the search for a compact representation of an author's writing style, or stylometric fingerprint. Applications of this study range from forensics to questions of plagiarism in the works of both published authors as well as students. With recent developments in computational efficiency and information processing, authorship attribution studies are of both increasing interest and accuracy. The study of authorship attribution, sometimes called stylometry, has its beginnings in works published over a century ago which proposed distinguishing authors by looking at word lengths and average sentence lengths. These two rudimentary ideas have improved since. A significant development came with the introduction of the influential idea of analyzing function words as a way to characterize authors' styles. Function words are words like prepositions, conjunctions, and pronouns which on their own carry little meaning but instead help define grammatical relationships between words. The study of function words is beneficial as they primarily inform about syntax rather than content. Since the introduction of function words as stylometric fingerprints many methods have been introduced to analyze the frequency of these words in texts written by different authors. Attention has also been given to analyzing features other than appearances of high-frequency words. Examples of these are the use of vocabulary richness, word stability – the extent to which a word can be replaced by an equivalent –, or syntactical markers like part-of-speech taggers.

Author	Texts	Author	Texts
Shakespeare, W.	10	Twain, M.	9
Austen, J.	7	Allen, G.	7
Cooper, J.F.	6	Dickens, C.	6
Marlowe, C.	8	Bacon, F.	6
Beaumont, F. & Fletcher, J.	7	Hawthorne, N.	6
Abbott, J.	7	James, H.	8
Alger, H.	7	Jonson, B.	6
Alcott, L.M.	7	Aldrich, T.B.	7
Garland, H	8	Melville, H.	8

Figure 18: List of authors used to test accuracy of classification with word adjacency networks. The total number of texts available for each author are also shown.

Our approach to authorship attribution focus on function words but instead of using their frequency distribution as an author signature we propose the use of the relational structure of function words. In order to classify the authorship of a text we compute an asymmetric network of function word adjacencies capturing how likely it is to find a particular function word within the next few words conditional on the occurrence of another given word. The resulting matrices can be interpreted as transition probabilities of a Markov chain. The similarity of different texts is then estimated by the relative entropy of these transition probabilities. We have tested the proposed methodology in authorship attribution problems including texts from up to 18 different authors using training sets consisting of between 1 and 6 known texts per author – see Fig. 19. Estimation accuracy in the order of at least 90% is observed in most cases. We have further demonstrated that our classifier performs better that classifiers based in word frequencies. Perhaps more important, numerical experiments show that classifiers based on word frequencies encode different stylometric fingerprints than the classifiers proposed here and can then be combined for increased attribution correctness.

	Known texts per author						Rnd.
Number of Authors	1	2	3	4	5	6	Attr.
2	1.00	1.00	1.00	1.00	1.00	1.00	.50
3	.87	1.00	1.00	1.00	1.00	1.00	.33
4	.76	.92	1.00	1.00	1.00	1.00	.25
5	.74	.90	1.00	1.00	1.00	1.00	.20
6	.74	.82	.93	.90	.93	1.00	.17
7	.63	.85	.94	.92	.94	1.00	.14
8	.67	.86	.94	.93	.95	1.00	.13
9	.60	.83	.92	.90	.95	.92	.11
10	.55	.77	.90	.91	.95	.92	.10
11	.53	.74	.89	.86	.88	.85	.09
12	.57	.76	.90	.87	.89	.87	.08
13	.60	.78	.91	.88	.90	.88	.08
14	.59	.78	.90	.86	.87	.88	.07
15	.61	.77	.89	.87	.88	.88	.07
16	.57	.73	.85	.84	.88	.89	.06
17	.58	.74	.85	.85	.89	.90	.06
18	.54	.69	.79	.83	.88	.86	.06

Figure 19: Accuracy for different number of candidate authors and number of known texts per author. Expected accuracy of random attribution is also informed. Accuracy decreases with increasing number of authors and decreasing number of training texts per author but remains large in most situations.

G.1 Ongoing work, acknowledgments and references

We are currently performing a comparative study of word adjacency networks for English authors of the Renaissance. From this study it can be easily seen that Bacon as well as Marlowe have stylometric fingerprints very different from the stylometric fingerprint of Shakespeare. This is, however, still on a preliminary stage. The technical ideas to generate and compare word adjacency networks for different authors come from the prolific imagination of Santiago Segarra. The actual legwork of comparing networks and running numerical analyses has been undertaken by Mark Eisen. Our preliminary results have appeared here:

[1] S. Segarra, M. Eisen, and A. Ribeiro, "Authorship attribution using function words adjacency networks," in *Proc. Int. Conf. Acoustics Speech Signal Process.*, vol. (submitted), Vancouver Canada, May 26-31 2013.