Networked Decision Systems

Munther A. Dahleh and Michael Rinehart

Decision and Communication Networks: Overview and Challenges

A *decision network* can be broadly characterized as a distributed system of locally controlled agents whose dynamics and/or objective functions have a neighborhood structure that can be described by a graph. The decision network is supported by an underlying *communication network* that may consist of both wired and wireless networks of varying quality and whose connectivity structure need not align with the decision network topology. We refer to the combination of the two networks as a *networked decision system*. A schematic networked decision system is shown in Fig. 1.

A familiar example of a networked decision system is a formation of unmanned aerial vehicles (UAVs). Each UAV has a local controller to control its flight, but it must also follow commanded trajectories while avoiding collisions and the like. This may require information from other nearby UAVs, ground bases, or other information sources. In addition, a leader UAV may need to provide trajectory or waypoint commands to the formation. These decisions can be communicated through the formation itself (as a multihop routing network) or through other nodes. Other examples of networked decision systems include distributed emergency response systems, interconnected transportation, energy systems, and even social networks.

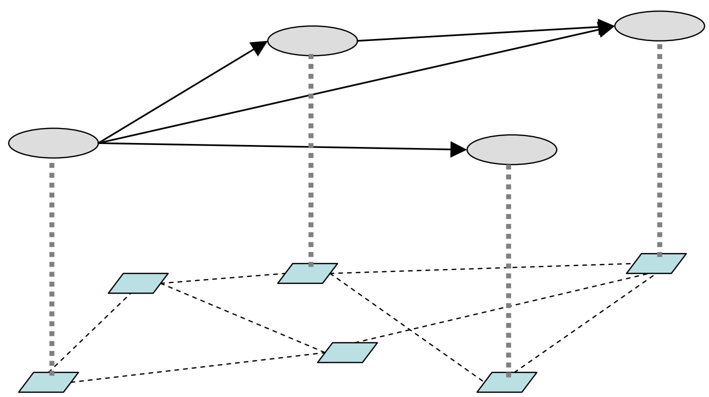

Figure 1. Illustration of a networked decision system. The upper-level nodes represent the decision network component (such as UAVs) and the lower-level nodes represent the communication network component (such as a multihop network).

Examples of networked decision systems include UAV formations, distributed emergency response systems, interconnected transportation, energy systems, and even social networks.

Networked decision systems are pervasive, and society and industry are becoming increasingly dependent on them. However, decentralized decision making over imperfect networks is fraught with difficulties. Issues and challenges are especially pronounced when dynamics are involved as the stability of the network also becomes a top priority. It is precisely these areas that the controls research community, with its history of designing robust and optimal dynamic systems, can address.

The ultimate objective of controls-related research in networked decision systems is a general analysis framework that can be used to derive fundamental performance limitations. The variety of realistic complications that such a framework must accommodate—communication delays, uncertainty in information, competitive environments, limitations of communication and computational resources, learning and adaptation, mobility in agents or infrastructure nodes—points to the ambitious nature of this goal. We begin our discussion with a description of the latest research in networked control systems. Although these efforts have revealed many fundamental limitations of these systems, generalizations of them lead to a general formulation of interest. We conclude with some broad considerations related to a unified theory of networked systems.

A research objective is to characterize the fundamental limitations and capabilities of networked systems by deriving performance bounds that are functions of the underlying topologies of the networks, the capacities of the communication links, the dynamics of each node, and the computational and storage resources available to each node.

Decisions Networks: Fundamental Limits and Open Questions

Challenges in Networked Decision Systems

If one is able (willing) to assume a priori that the decision network and communication network do not interact except though an interface of constraints and requirements, the two networks can be analyzed and designed essentially independent of one another, allowing for a classical analysis of the system. In particular, the communication network can be abstracted as a set of static constraints (such as channel capacities or delays) on the operations of the decision network, whereas the decision network can be seen by the communication network as imposing requirements or preferences such as performance guarantees or utility functions. However, this assumption is rarely true in practice. For example, the decision network may take actions that disconnect the communication network or the communication network may not efficiently route critical information to portions of the decision network quickly enough, affecting performance or even stability.

Complicating matters are the dynamics of the decentralized decision network itself. Even if the underlying communication network is perfect (infinite capacity and no latency), the decision network possesses performance limitations that are missing in the centralized single decision agent. In fact, the analysis and design of distributed systems with different information patterns is still an open problem. Resource constraints that necessitate practical protocols and algorithms, and even fundamental challenges in control theory such as delays, further complicate this setting.

Control theory, information theory, optimization and game theory, and graph theory considered aspects of these applications in isolation and were able to provide basic limitations such as those captured by Bode's integral formula, Shannon's information transmission, Myerson-Satterthwaite's result on bilateral trade, and the spectral theory of graphs. However, no general analysis framework exists that is capable of addressing the interplay of these factors. In fact, the very paradigms for control and communication systems are incompatible. For example, although information theory has focused on zero-error transmission with possibly large delays, control systems tend to be very sensitive to delays while being less sensitive to static and dynamic errors—both consequences of the use of feedback. Even in the context of a single agent, the interplay between the physical space (where the agent is expected to perform) and the information space (which is described by the ability to communicate) when the agent has limited resources and when the success of communication depends on the actual dynamical behavior is still to be investigated. A network of such cooperating agents creates an even more challenging set of problems in terms of the fundamental limits. Finally, a network of autonomous agents that have possibly conflicting interests is still more difficult to analyze and coordinate.

The results presented in this section provide insight into the fundamental performance limits of decision networks. Some of these limitations arise from dynamics of the decentralized nature of the system and others from the agent's uncertainties about the system, in part due to delay and channel rate limitations.

Single Agent: The Value of Side Information in Static Decisions

The network can often be part of the overall design of a system. In complex applications such as transportation systems or the power grid, which involve humans in the loop, it is critical that only select information be communicated to the decision maker. Otherwise decisions will be delayed substantially until the decision maker sorts through the massive data sets. In short, information must be compressed and filtered so that only the information that most influences decision making is communicated. Also, because gathering, transmitting, and storing information can be costly, the minimum amount of information that is required to reach a certain level of performance should be determined.

To begin to understand the relationship between information type, information quantity, and decision making, consider a simple but prototypical problem: a single agent traversing the shortest path of a graph [1], [2]. Although only a single agent is considered, the information on which this agent bases its decisions—uncertain, intermittent, partial information about traversal delays along different edges—is analogous to the problems that would be faced if such information were being communicated by distributed sensors over an imperfect communication network. Information-theoretic bounds and other results from the single-agent scenario carry over to networked decision systems.

The standard stochastic shortest-path problem can be described briefly as follows. An agent wishes to traverse a graph along the shortest path in that graph. The delays on the edges are random (with a known distribution), and the agent may or may not know some information about the edge delays in advance of choosing a path. Now if the agent has limited resources with which to gather information about the edge delays in advance of its travel (for example, it has a limited budget for purchasing sensors), relevant questions in this context are: What types of sensors should the agent purchase, and on which edges should the agent place the sensors to best improve its overall performance? Beyond shortest-path optimization, we may more generally seek to provide a simple, intuitive framework for studying decision making under limited information conditions as well as to provide algorithms that (sub)optimally allocate information resources (such as sensors or bandwidth) to best improve the agent's performance.

In this static, centralized, and performance-centric setting, the optimal information is not characterized by mutual information quantities or bit rates, but rather as a measure of the degree to which that information is concentrated to the agent's decision subspace (termed the *actionable information*). In particular, the agent uses the information to estimate the edge delays, and the variance of this estimate in a particular subspace is the sole determinant of the agent's performance [1]. In fact, the agent's overall estimation error is irrelevant and can be arbitrarily high. Furthermore, under certain conditions, a practical scheme exists by which the agent can guarantee that the information it receives is concentrated (that is, without additional processing) to its actionable component: place all sensors to at most two paths of the graph. Generally, this scheme may contain some irrelevant information, but the performance resulting from this configuration can be shown to be acceptable.

This setting can be further generalized to a quasi-dynamic setting [2] where information is gradually revealed to the agent as it traverses the graph. In this case, the actionable information changes with each step the agent takes. If future information is reconcentrated to these subspaces, the agent's performance can be shown to further improve. However, if the information is blindly broadcast, the agent's performance can only degrade.

Designing a network with limited capacity to support decision making brings to light important research questions that generalize this framework:

- Inclusion of dynamics: The amount of actionable information determines the agent's performance. How can this notion be generalized in a non-performance-centric setting where the agent has dynamics and is concerned with stability?

- Algorithms for computing the actionable information set: The actionable information was shown to correspond to a subspace in the single-agent shortest-path problem. In more general settings with nonlinear objectives and multiple agents, the actionable information set may not have such a simple characterization. Can general techniques be developed for efficiently approximating the set of actionable information in this case?

Single Agent: Stability and Asymptotic Performance Under Communication Constraints

Understanding the fundamental limitations of performance in a feedback system is critical for effective control design. Substantial progress has been made in this direction, addressing questions of stability and performance tradeoffs in feedback systems. One of the most powerful results capturing performance tradeoffs in a stable linear feedback system is Bode's integral formula [3], which captures performance limitations in terms of the unstable modes of the plant.

In the context of centralized control under communication constraints, generalizations to this result as well as other results were obtained using information-theoretic concepts. For example, research has shown that the minimum bit rate through a discrete, error-free channel between the plant and controller that is required to stabilize a linear system is expressed purely in terms of the unstable modes of the plant [4], [5]. Furthermore, practical communication schemes can be developed that provide that base rate. A performance-centric variation of this problem is considered where the plant and controller have perfect communication but track a reference that is communicated over a channel [6]. Furthermore, the controller is to provide good model-matching performance subject to this limited reference. Research shows that there is an inherent tradeoff between communication delay and performance which forces the design of the encoder/decoder and the controller to be performed simultaneously [6].

In the two cases above, communication constraints were treated as bit-rate constraints on a discrete channel. A different representation for communication constraints is considered whereby a communication channel between the plant and controller is characterized solely by its capacity [7], [8]. A nonclassical analysis using information-theoretic quantities is used to examine the flow on entropy in the feedback loop as a means of obtaining fundamental asymptotic performance limitations. The result is a generalization of Bode's integral formula that provides conditions under which this limit can be improved by using side information.

To apply an entropy-flow analysis, properties of the controller must be characterized in terms of information-theoretic constraints. The causality of the controller and overall stability of the plant are expressed, respectively, in terms of a mutual information equality and a variance constraint [7], [8].

Generally, such abstract representations of the system allow for an asymptotic analysis that can reveal fundamental performance limits.

Although these results bridge the gap between control and communication, much remains to be explored. Following are some interesting open problems:

- Notions for information: What is the correct notion of "information" when communication supports a decision system? The notion of information captured by Shannon in point-to-point communication is not adequate in this setting. In the context of channel coding, block codes perform optimally in transmitting a message with small probability of error; however, such codes can be detrimental to a control system due to large delays.

- Tradeoff between bit rate and delay: How do we address the interplay between control and communication? The summary above assumes that the system dynamics are decoupled from the communication channel. In many situations, the bandwidth or capacity of the channel depends directly on the state of the underlying dynamic systems, such as in a mobile system where the communication depends on its actual physical location. Since the mobile system can choose to deploy itself at a particular location, the power consumed is shared with the power available for communication. Such examples where communication directly interferes with the control strategy are not very well understood.

Network of Cooperating Agents: Decentralized Computation Under Communication Constraints

We now move beyond the centralized decision maker setting to a decentralized setting, specifically decentralized decisions over unreliable networks. Examples of such networks include ad hoc wireless networks, satellite networks, and noisy social and human networks. Such networks can severely limit the capabilities of decision makers as their ability to estimate the underlying states of the systems is limited by the ability to faithfully communicate with the other agents in a timely fashion. The research objective is to characterize the fundamental limitations and capabilities of such networked systems by deriving performance bounds that are functions of the underlying topologies of the networks, the capacities of the communication links, the dynamics of each node, and the computational and storage resources available to each node.

When nodes can have unlimited computational power, research has shown that the conductance of the network graph—a measure of how "well knit" the graph is—plays a critical role in characterizing the performance of consensus-type problems where nodes are trying to compute a function of a set of initial values that are distributed over the network [9]. In particular, the time needed for each node to compute an accurate estimate of its function scales as the inverse of the conductance. For example, a ring network that communicates with neighbors with probability 1/4 scales as the inverse of the number of vertices, which implies a linear growth in convergence time for the estimates. Networks that communicate with all agents with the same probability have no bottlenecks and their conductance is constant regardless of the network's size. For example, the preferential model of the Internet has this property, which indicates that the Internet is a good medium for distributed computation. Another example of such a network is the ad hoc wireless model of Gupta-Kumar [10], which allows two wireless devices to communicate simultaneously only if they are outside a disk of a certain radius (this is often referred to as the disk model). In this case, the computation can be obtained accurately at a rate not faster than the square root of the number of vertices.

A natural generalization of this framework is one where evolving functions need to be communicated. This problem is further complicated in the realistic case of agents having dynamics. For example, if agents communicate with other agents over channels with capacities that depend on their locations and resources, the graph connecting them may change dynamically. Putting aside the agent's own dynamics, the stability of the distributed function estimation itself is put in danger as the graph changes.

Previous work has also explored some of the mechanisms for computation in the presence of varying time delays and changes in network connectivity [11], [12], but only relatively simple operations such as consensus protocols have been fully explored. Conversely, some work has been done in maintaining robust communications topologies, but without regard for the most effective utilization of network resources or the details of the desired information flow and possible effects of latency. These problems are particularly difficult in the case where local decisions are made at the network's nodes, requiring global properties to either be represented in a distributed fashion or estimated by individual nodes (including receivers and transmitters in the network).

In addition to the above, further research areas include:

- Architectural limitations on distributed problems: Consider, for example, a network where agents can only communicate their decisions (or the values of the functions they are computing). In this context, we think of these functions as utilities. Communicating utilities gives only aggregate information about the underlying state of the system and imposes severe limitations on the ability to learn the state. How can these limitations be characterized?

- Robustness: In this regard, it may be beneficial to search for the right topology (or metric) on the set of graphs that is amenable to perturbation analysis. Under what perturbation conditions is asymptotic estimation possible?

Network of Competitive Agents: Information Aggregation and Asymptotic Learning

Social networks are attracting substantial attention within the research community. In particular, a tremendous opportunity exists for bringing in quantitative tools to analyze the formation of such networks as well as to study the impact of such networks on decision making. What differentiates such networks from standard decentralized networks is the human presence. A question that arises in the investigation of networks with human actors is how game-theoretic interactions modify the well-known existing results on dynamic aggregation of decentralized information over networks with non-autonomous agents (for example, see the literature on consensus [13]-[17]).

Results have been reported that begin to address this framework [19]. They show that when selfish agents are sequentially detecting an underlying binary state of the world, information may not aggregate properly. The loss of collective wisdom is due to the "herding" phenomenon often witnessed in technology and fads. A realistic framework for learning in a multi-agent system must model the structure of social networks with which individuals observe and communicate with each other; however, such generalizations turn out to be challenging to analyze. One difficulty with this class of models is that to determine how beliefs will evolve, we need to characterize the perfect Bayesian Nash equilibrium, which involves rather complex inferences by individuals. To this end, we will consider a simplified model where the agents observe the actions of a neighborhood of individuals that are randomly chosen from the entire set of past actions of this neighborhood. Although actual social leaning can involve very complicated dynamics not captured in this simplified model, it does provide a first-order approximation for which definitive statements can be made, and the fundamental limitations of this model may hold in more complicated models.

More recent work addressed this problem and established exact conditions under which herding is impossible [20]. In this work, these effects are captured in terms of characteristics of the graph's interconnectedness and the properties of the underlying random process. In particular, under certain conditions, an excessively influential group can emerge within a social network if interconnectedness among individuals is not rich enough.

These results consider an idealized situation where all agents have the same utility and where there is an absence of disruption. Furthermore, they only address asymptotic learning as the size of the network increases. Hence, several interesting research directions in this field have not been pursued or have provided only partial results:

- Sequential decisions and feedback: Analysis was simplified by allowing agents to fix a decision once it is made, but repeated decision making better reflects real-world dynamics. How do repeated decisions and endogenous sequencing of actions affect asymptotic learning over time?
- Perturbations: The influence of external effects (such as media, injecting outside agents, changing the network topology) on the propagation of beliefs is relevant because such outside effects can serve as either control inputs or as adversarial influences on the system. For example, what types of networks allow a "reversal" in the beliefs of individuals?
- Forward-thinking agents: In social learning models studied to date, individuals care only about their immediate payoffs. What general approach should one take toward analyzing the perfect Bayesian Nash equilibrium in the case where agents' payoffs depend on the future decisions of other agents?

Broad Considerations in Decision and Communication Networks

The natural generalizations considered for each of the previous works seem to quickly lead to common problems of high importance. Although the specifics of these problems still vary (each has different objectives and algorithms), a general analysis framework could be established that can be used to derive fundamental performance limitations. Below we discuss several research areas that may be helpful in developing such a framework.

- Network separation principle: The separation principle from classical control theory offers conditions under which a feedback control signal cannot improve the controller's estimate of the plant's state. When it applies, the optimal performance of the system can be directly analyzed by constructing an optimal estimator and controller. However, if these conditions are not met, even a simple feedback system can have a complex optimal controller. Under what conditions are system uncertainties independent of the agents' decisions? A "degree of separation" may be useful in establishing approximate results in this challenging area.
- Dynamic notions for actionable information: In learning and centralized-feedback control, an entropy-flow analysis was used to study the dynamic exchange of information between agents. The results were algorithmically free, asymptotic fundamental limits for performance. However, in the performance-centric setting of shortest-path optimization, it was the amount of information concentrated to the actionable subspace of decisions that affected the quality of

decision making. In a dynamic setting, does there exist a similar set to which information should be concentrated and that changes over time? How should agents track it? The flow of actionable information content over the network may yield tighter, more useful fundamental limits than entropy flow alone.

- Representations for abstract computation: An effort to link decision making to information flow may require developing a representation for algorithms in the language of information flow. In the case of a centralized-controller feedback system, causal, stabilizing controllers can be represented by imposing information-theoretic constraints on the feedback system. Can such formulations be extended to decentralized and nonlinear settings? Additional constraints that may be useful to develop are those that capture limited computational capability.

- Representations for communication: The notion of information captured by Shannon in point-to-point communication is not adequate for analysis. As noted earlier, although block codes perform optimally in the context of transmitting a message with small probability of error, such codes can be detrimental to a control system due to large delays. How can we efficiently represent causality across a network with many information flows? Notions of mutual information and information rates do not completely capture the interactions of multiple causal dependences.

- Robustness to perturbations: Perturbations in network topology, computation, or communication may propagate errors throughout the network that can degrade performance or, worse, result in positive feedback loops in the system that may amplify the effect of the errors, destabilizing the system. To illustrate the types of perturbations that need to be specially considered in dynamic agent networks, consider the case where the interaction between the agent's dynamics and the graph are carefully designed, but a time-varying perturbation in the graph results in a transient cycle in information flow. If the network is a Bayesian learning network, these cycles may destabilize learning. Even in the simplest case where the dynamics of the nodes can be modeled as linear input/output systems (including time delays), the static graph structure is known to be crucial for determining its overall stability [21].

Selected recommendations for research in networked decision systems:

- What is the correct notion of "information" when communication supports a decision system, and how do we address the interplay between control and communication? Fundamental problems of analysis and design in cases where communication directly interacts with the control strategy need to be investigated.
- Our understanding of the fundamental limitations and capabilities of decentralized networked systems under uncertainties is incomplete. Performance bounds that are functions of the underlying topologies of the networks, the capacities of the communication links, the dynamics of each node, and the computational and storage resources available to each node would be useful for many applications.
- Connections with game theory are an important research area, with several open problems. For example, how do repeated decisions and endogenous sequencing of actions affect asymptotic learning over time in a game-theoretic network of competitive agents?

References

[1] M. Rinehart and M.A. Dahleh. "The value of side information in shortest path optimization," to appear in *IEEE Transactions on Automatic Control,* 2010.

[2] M. Rinehart and M.A. Dahleh. "The value of sequential information in shortest path optimization," submitted to *Automatica,* 2010.

[3] H.W. Bode. *Network Analysis and Feedback Amplifier Design.* Princeton, NJ: D. Van Nostrand, 1945.

[4] S. Tatikonda and S. Mitter. "Control over noisy channels," *IEEE Transactions on Automatic Control,* vol. 49, pp. 1196-1201, July 2004.

[5] S. Tatikonda and S. Mitter. "Control under communication constraints," *IEEE Transactions on Automatic Control,* vol. 49, pp. 1056-1068, July 2004.

[6] S. Sarma and M.A. Dahleh. "Remote control over noisy communication channels: A first-order example," *IEEE Transactions on Automatic Control,* vol. 52, pp. 284-289, February 2007.

[7] N.C. Martins and M.A. Dahleh. "Feedback control in the presence of noisy channels: Bode-like fundamental limitation of performance," *IEEE Transactions on Automatic Control*, vol. 53, pp.1604-1615, August 2008.

[8] N.C. Martins, M.A. Dahleh, and J.C. Doyle. "Fundamental limitations of disturbance attenuation with side information," *IEEE Transactions on Automatic Control,* vol. 52, pp.56-66, January 2007.

[9] O. Ayaso, D. Shah, and M.A. Dahleh. "Information theoretic bounds for distributed computation," submitted to *IEEE Transactions on Information Theory,* 2008.

[10] P. Gupta and P.R. Kumar. "The capacity of wireless networks," *IEEE Transactions on Information Theory,* vol. 46, pp. 308-404, March 2000.

[11] A. Jadbabaie, J. Lin, and A.S. Morse. "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control,* vol. 48, no. 6, pp. 988-1001, June 2003.

[12] R. Olfati-Saber and R. M. Murray. "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control,* vol. 49, pp. 1520-1533, September 2004.

[13] J.N. Tsitsiklis. "Problems in Decentralized Decision Making and Computation," PhD thesis, Massachusetts Institute of Technology, 1984.

[14] J.N. Tsitsiklis, D.P. Bertsekas, and M. Athans. "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Transactions on Automatic Control,* vol. 31, no. 9, pp. 803-812, September 1986.

[16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. "Gossip algorithms: Design, analysis, and applications," in Proceedings of IEEE INFOCOM, vol. 3, pp. 1653-1644, March 2005.

[17] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. "Convergence in multiagent coordination, consensus, and flocking," in Proceedings of IEEE CDC, pp. 2996-3000, December 2005.

[18] A. Nedić and A. Ozdaglar. "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control,* vol. 54, pp. 48-61, January 2009.

[19] A. Nedić, A. Olshevsky, A. Ozdaglar, and J.N. Tsitsiklis. "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control,* vol. 54, pp. 2506-2517, November 2009.

[20] S. Bikchandani, D. Hirshleifer, and I. Welch. “A theory of fads, fashion, custom, and cultural change as information cascades,” *Journal of Political Economy,* vol. 100, pp. 992-1026, 1992.

[21] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian Learning in Social Networks, LIDS Technical Report 2779, 2008.

[22] J.A. Fax and R.M. Murray. “Information flow and cooperative control of vehicle formations,” *IEEE Transactions on Automatic Control,* vol. 49, no. 5, pp. 1465-1476, September 2004.