Chapter One

Introduction

> "If a man writes a book,
> let him set down only what he knows.
> I have guesses enough of my own." — Goethe

In this introductory chapter, we provide a brief discussion of networked multiagent systems and their importance in a number of scientific and engineering disciplines. We particularly focus on some of the theoretical challenges for designing, analyzing, and controlling multiagent robotic systems by focusing on the constraints induced by the geometric and combinatorial characters of the information-exchange mechanism.

1.1 HELLO, NETWORKED WORLD

Network science has emerged as a powerful conceptual paradigm in science and engineering. Constructs and phenomena such as interconnected networks, random and small-world networks, and phase transition nowadays appear in a wide variety of research literature, ranging across social networks, statistical physics, sensor networks, economics, and of course multiagent coordination and control. The reason for this unprecedented attention to network science is twofold. On the one hand, in a number of disciplines–particularly in biological and material sciences–it has become vital to gain a deeper understanding of the role that inter-elemental interactions play in the collective functionality of multilayered systems. On the other hand, technological advances have facilitated an ability to synthesize networked engineering systems–such as those found in multivehicle systems, sensor networks, and nanostructures–that resemble, sometimes remotely, their natural counterparts in terms of their functional and operational complexity.

A basic premise in network science is that the structure and attributes of the network influence the dynamical properties exhibited at the system level. The implications and utility of adopting such a perspective for engineering networked systems, and specifically the system theoretic consequences of such a point of view, formed the impetus for much of this book.[1]

---

[1] One needs to add, however, that–judging by the vast apparatus of social networking, e.g.,

1.2 MULTIAGENT SYSTEMS

Engineered, distributed multiagent networks, such as distributed robots and mobile sensor networks, have posed a number of challenges in terms of their system theoretic analysis and synthesis. Agents in such networks are required to operate in concert with each other in order to achieve system-level objectives, while having access to limited computational resources and local communications and sensing capabilities. In this introductory chapter, we first discuss a few examples of such distributed and networked systems, such as multiple aerospace vehicles, sensor networks, and nanosystems. We then proceed to outline some of the insights that a *graph theoretic* approach to multiagent networks is expected to provide, before offering a preview of the book's content.

1.2.1 Boids Model

The Reynolds boids model, originally proposed in the context of computer graphics and animation, illustrates the basic premise behind a number of multiagent problems, in which a collection of mobile agents are to collectively solve a global task using local interaction rules. This model attempts to capture the way social animals and birds align themselves in swarms, schools, flocks, and herds. In the boids flocking model, each "agent," in this case a computer animated construct, is designed to react to its neighboring flockmates, following an ad hoc protocol consisting of three rules operating at different spatial scales. These rules are *separation* (avoid colliding with neighbors), *alignment* (align velocity with neighbors' velocities), and *cohesion* (avoid becoming isolated from neighbors). A special case of the boids model is one in which all agents move at the same constant speed and update their headings according to a nearest neighbor rule for group level alignment and cohesion. It turns out that based on such local interaction rules alone, velocity alignment and other types of flocking behaviors can be obtained. An example of the resulting behavior is shown in Figure 1.1.

1.2.2 Formation Flight

Distributed aerospace systems, such as multiple spacecraft, fleets of autonomous rovers, and formations of unmanned aerial vehicles, have been identified as a new paradigm for a wide array of applications. It is envisioned that distributed aerospace technologies will enable the implementation of a spatially distributed network of vehicles that collaborate toward

---

email, facebook, twitter, and a multitude of networked, coordinated, and harmonic behavior in nature and the arts–our fascination with multiagent networks is more intrinsic.

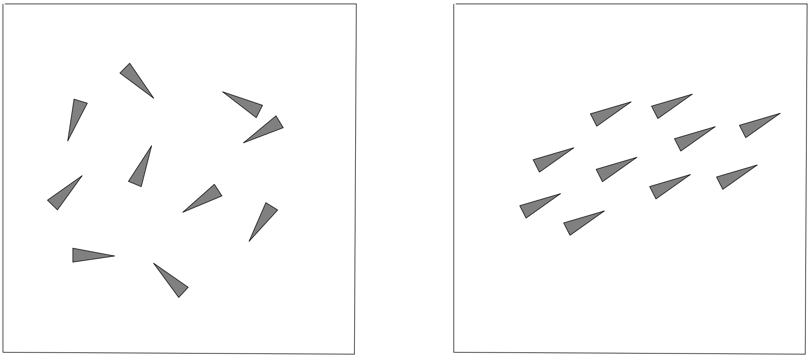

Figure 1.1: A Reynolds boids model in action. Ten agents, each with an arbitrary initial heading (given by the orientation of the triangles) and spacing, are considered (left); after a while they are aligned, moving in the same general direction at regular interagent distances (right). When this is the case, we say that *flocking* has been achieved.

a single collective scientific, military, or civilian goal. These systems are of great interest since their distributed architecture promises a significant cost reduction in their design, manufacturing, and operation. Moreover, distributed aerospace systems lead to higher degrees of scalability and adaptability in response to changes in the mission goals and system capabilities.

An example of a multiple platform aerospace system is space-borne optical interferometry. Space interferometers are distinguished by their composition and operational environment. They are composed of separated optical instruments, leading to a so-called sparse aperture. Although optical interferometers can, in principle, function on the earth's surface, there are many advantages in operating them in space. Space-borne interferometers have greater optical sensitivity and resolution, wider field of view, and greater detection capability. The resolution of these interferometers, as compared with space telescopes (e.g., Hubble), is dictated by the *separation* between the light collecting elements (called the baseline) rather than their *size*. Consequently, as the achievable imaging resolution of a space telescope is dictated by advanced manufacturing techniques, the size of the launch vehicle, and the complex deployment mechanism, the capability of a space-borne optical interferometer is limited by how accurately the operation of separated optical elements can be coordinated. These space-borne optical interferometers can be mounted on a single large space structure, composed of rigid or semirigid trusses or even inflatable membranes. In this case, the structural dynamics of the spacecraft plays a major role in the operation and the

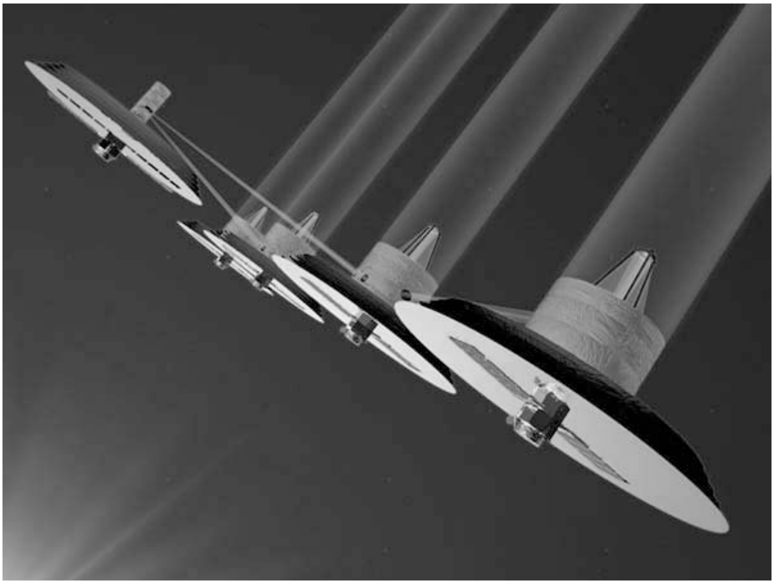

Figure 1.2: Terrestrial Planet Finder, courtesy of JPL/NASA

success of the mission. An alternate approach is to fly the interferometer on multiple physically separated spacecraft, that is, a distributed space system. An example of such a mission is the Terrestrial Planet Finder (TPF) shown in Figure 1.2.

Another important set of applications of networked aerospace systems is found in the area of unmanned aerial vehicles of various scales and capabilities. These vehicle systems provide unique capabilities for a number of mission objectives, including surveillance, synthetic aperture imaging, mapping, target detection, and environmental monitoring.

1.2.3 Sensor Networks

A wireless sensor network consists of spatially distributed autonomous devices that cooperatively monitor physical or environmental conditions, such as temperature, sound, vibration, or pressure. Each node in a sensor network is equipped with a wireless communication device as well as an energy source–such as a battery–that needs to be efficiently utilized. The size, cost, and fidelity of a single sensor node can vary greatly, often in direct correspondence with its energy use, computational speed, and the ease by which it can be integrated within the network. Each sensor exchanges information on its local measurements with other nodes in the network in order to reach an accurate estimate of the physical or environmental variable of interest. We note that the efficiency requirement on the utilization of the energy source for each sensor often dictates a geometry on the internode communication for the sensor network.

1.2.4 Nanosystems

Recently, there has been a surge of interest by material scientists in organic compounds that are interconvertible via chemical reactions; this process is often referred to as *tautomerization*. These chemical reactions can be used for constructing molecular switches, where a molecule is steered between two or more stable states in a controlled fashion. Other electronic components such as diodes and transistors can be made that rely on similar induced transitions between structural isomers. Such molecular devices can then be put together, leading to the possibility of designing molecular circuits, networks, and more generally, *molecular dynamic systems*. An example of a molecular switch is a hydrogen tautomerization employed to manipulate and probe a naphthalocyanine molecule via low-temperature scanning tunneling microscopy. The properties and functionality of the corresponding molecular machines and networks are highly dependent on the inter-molecular bonds that can generally be manipulated by techniques such as electron beam lithography and molecular beam epitaxy.

1.2.5 Social Networks

Social networks are comprised of social entities, such as individuals and organizations, with a given set of interdependencies. The interaction between these entities can assume a multitude of relations, such as financial, social, and informational. Such networks are of great interest in a variety of fields, including theoretical sociology, organizational studies, and sociolinguistics. In fact, the *structure* of social networks has always been of fundamental importance for understanding these networks. More recently, the notion of manipulating the network structure has been contemplated as a viable means of altering the network behavior. For example, the concept of a *change agent* refers to a network entity that intentionally or indirectly causes or accelerates social, cultural, or behavioral change in the network.

1.2.6 Energy Networks

Complex, large-scale energy systems, delivering electrical and mechanical energy from generators to loads via an intricate distribution network, are among the most useful engineered networked dynamic systems. These systems often consist of a heterogeneous set of dynamic systems, such as power electronics and switching logics, that evolve over multiple timescales. Dynamics, stability, and control of individual power system elements (e.g., synchronous machines) or their interconnections (e.g., multi-machine models) have extensively been examined in the literature. However, as the need for more efficient generation and utilization of energy has become prevalent, distributed and network architectures such as the “smart grid” have gained particular prominence.

1.2.7 The Common Thread

The examples above, sampled from distinct disciplines, share a set of fundamental system theoretic attributes with a host of other networked multiagent systems. In a nutshell, such systems consist of (1) dynamic units, potentially with a decision making capability and means by which they can receive and transmit information among themselves, and (2) a signal exchange network, which can be realized via wired or wireless protocols in engineering, biochemical reactions in biological systems, and psychological and sociological interactions in the context of social networks.

The fundamental feature of networked systems, distinguishing them from systems that have traditionally been considered in system theory, is the presence of the network and its influence on the behavior of the overall system. Consequently, a successful “system theory for networked systems” has to blend the mathematics of information networks with paradigms that are at the core of dynamic system theory (stability, controllability, optimality, etc.). One of the challenging aspects facing such an interdisciplinary marriage in the context of system theory is that many network properties, for example, the network geometry, have a logical or combinatorial character.

1.3 INFORMATION EXCHANGE VIA LOCAL INTERACTIONS

In order to have a concrete model of “local interactions,” in this section, we delineate the local nature of information exchange mechanisms for robotic networks.

1.3.1 Locality in Communication

One way in which agents can share information with their surroundings is through communication channels. But transmitting and receiving information requires energy, which is typically a sparse commodity in many networked applications, such as sensor networks and mobile ad hoc communication networks. Hence, only agents within a limited communication range can exchange information directly, forcing information to propagate through the network over intermediary nodes. Another communication constraint pertains to the available bandwidth. If a large collection of agents simultaneously broadcast large amounts of data, the communication channels saturate and lead to sharp deterioration of the communication system. Thus, in large networks, the information exchange should be maintained and kept parsimonious in order to satisfy the bandwidth limitations.

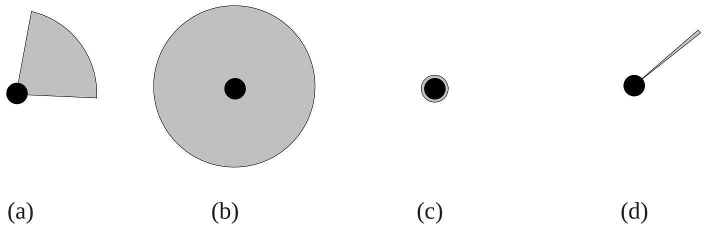

Figure 1.3: Various sensing geometries: (a) a vision-sensor with a wedge-shaped effective geometry; (b) an omnidirectional range sensor with a limited sensor range; (c) a tactile sensor provides information about the immediate surroundings; and (d) a single ray range sensor.

1.3.2 Locality in Sensing

Direct communications aside, agents can also infer information about each other and their environment through sensors. But every sensor has its own limitations in terms of range and resolution. Some of the most common sensors and their corresponding constraints are:

- **Vision-based sensors:** Cameras typically have long effective range (at least monocular vision), but they cover only a particular wedge-shaped region, as seen in Figure 1.3(a).

- **Range sensors:** The most common range sensors include sonars, infrared sensors, and laser scanners. These range sensors have very different sensing resolutions, ranging from very short range (e.g., low-cost infrared sensors) to covering hundreds of meters (e.g., high-quality laser-scanners). These sensors emit rays along a single direction but are typically ring-mounted to provide omnidirectional sensing capabilities (e.g., sonar and infrared rings) or have moving mirrors to provide scans across a larger area (e.g., laser scanners). This is shown in Figure 1.3(b).

A number of other sensing modalities are also widely used, with their own geometric constraints, as seen in Figure 1.3(c - d). However, as will be discussed throughout this book, this geometry will be subsumed by a graph theoretic interpretation of interactions as edges in the so-called *proximity graphs*, in which the existence of an edge indicates that neighboring nodes are within sensing range of each other.

1.4 GRAPH-BASED INTERACTION MODELS

The interaction geometry will indeed play an important role in the analysis and synthesis of networked multiagent systems regardless of whether the information exchange takes place over a communication network or through active sensing, or for that matter whether it assumes a wireless, chemical, physical, or sociological character. It turns out, however, that making the interaction protocol and its geometry explicit in the system-level analysis and control synthesis is far from trivial. In this direction, it becomes judicious to treat interactions as essentially combinatorial–at least initially–to codify whether an interaction exists and to what degree. An example of this abstraction is seen in Figure 1.4, in which the interaction geometry is defined by omnidirectional range sensors. As we will see throughout this book, such an abstraction, which cuts through the particular realization of the interaction, allows us to highlight the role of the interconnection topology, not only in the analysis of these systems but also in their synthesis.

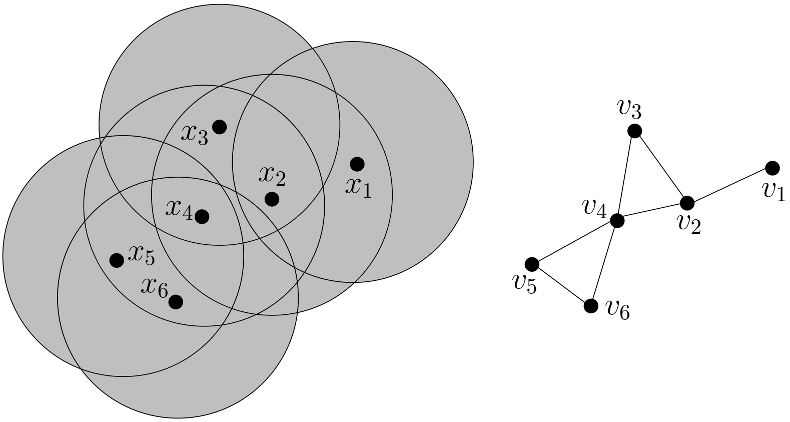

Figure 1.4: A network of agents equipped with omnidirectional range sensors can be viewed as a graph, with nodes corresponding to the agents and edges to the interactions.

1.4.1 Static, Dynamic, and Random Networks

If the edges in graphs are to be interpreted as enabling information to flow between the vertices on the corresponding edge, these flows can be directed as well as undirected. In other words, it is possible that the information will flow only in one direction. This would, for example, be the case if the vertices correspond to sensor agents, and agent i can sense agent j, while agent j can not sense agent i, for instance, due to different sensing modalities. In that case, the edge would be directed, with v_j as its "tail" and v_i as its "head." We will pictorially depict this as an arrow originating from v_j and ending at v_i. If the edge is undirected, we will simply drop the arrow and draw the edge as a line between the vertices.

However, directionality is not the only aspect of the edges that we will consider. We will also investigate different forms of temporal persistence, that is, situations in which the edges may disappear and reappear. In particular, we will group graphs into three classes:

- **Static Networks:** In these networks, the edges are static, that is, the edge set will not be time varying. This is, for example, the situation when a static communication network has been established, through which the information is flowing.

- **Dynamic, State-dependent Networks:** Here the edge set is time varying in that edges may disappear and reappear as functions of the underlying state of the network agents. For example, if the vertices in the graph correspond to mobile robots equipped with range sensors, edges will appear as agents get within the sensory range of each other, and be lost as agents get out of the sensory range.

- **Random Networks:** These networks constitute a special class of dynamic networks in that the existence of a particular edge is given by a probability distribution rather than some deterministic, geometric sensing condition. Such networks arise, for example, in the communications setting when the quality of the communication channels can be modeled as being probabilistic in nature.

It should be noted already at this point that these three types of networks will require different tools for their analysis. For static networks, we will rely heavily on the theory of linear, time-invariant systems. When the networks are dynamic, we have to move into the domain of hybrid systems, which will inevitably lead down the path of employing Lyapunov-based machinery for switched and hybrid systems. The random networks will in turn rely on a mix of Lyapunov theory and notions from stochastic stability.

1.5 LOOKING AHEAD

Graphs are inherently combinatorial objects, with the beauty but also limitations that come with such objects. Even though we will repeatedly connect with combinatorics, a host of issues pertaining to multiagent networks do not fruitfully lend themselves to a (pure) graph theoretic paradigm–at least not yet! Examples of such application domains include coverage control in sensor networks, which involves explicit partitioning of the environment and feedback control over a lossy and delayed network, where issues of delays, packet loss, and asynchronous operation, even for a pair of agents, are dominant. Moreover, the perspective adopted in this book does not include a detailed analysis of the underlying communication protocols, but instead employs a rather idealized model of information sharing, such as broadcast or single- and multi-hop strategies, and it is assumed that we can transmit and receive real numbers rather than quantized, finite bandwidth packets.

Another broad approach that we have adopted in this book is to work for the most part with simplified dynamics for the agents, that is, those with single and double integrators, linear time-invariant models, and unicycle models. In contrast, real-world networked systems are often comprised of agents with nontrivial dynamic input-output characteristics, interacting with each other via an elaborate set of interaction protocols. In this case, the behavior of the overall system depends not only on the interconnection topology and its detailed attributes, but also on how the interconnection protocol combines with the nonlinear and hybrid nature of the agents' dynamics.

Examples of topics that will be examined in this book include local interaction protocols for

- **Consensus:** having agents come to a global agreement on a state value;
- **Formations:** making the agents move to a desired geometric shape;
- **Assignments:** deciding a fair assignment of tasks among multiple agents;
- **Coverage:** producing maximally spread networks without making them disconnected or exhibit "holes" in their coverage;
- **Flocking/Swarming:** making the agents exhibit behaviors observed in nature, such as flocking birds, schooling fish, or swarming social insects;
- **Social Networks and Games:** analyzing how the outcomes of games and social interactions are influenced by the underlying interaction topology; and
- **Distributed Estimation:** organizing a group of sensors to collectively estimate a random phenomena of interest.

In later parts, we will also look at system theoretic models of controlled networks, capturing to what extend the behavior of networks can be influenced by exogenous inputs. We will examine dynamic notions of graph processes, thus allowing the graph structure itself be subject to control and time evolution. We conclude the book by providing an account of a framework for analyzing higher-dimensional interaction models via simplicial complexes.

NOTES AND REFERENCES

The boids model is due to Reynolds, who was motivated by animating movements of animal flocking [205]; this model was later employed by Vicsek, Czirók, Ben-Jacob, Cohen, and Shochet [238] for constant speed particles, mainly as a way to reason about self-organizing behaviors among large numbers of self-driven agents. This so-called Vicsek model, in turn, has provided an impetus for system theoretic analysis, such as the work of Jadbabaie, Lin, and Morse [124], which is also related to works on parallel and distributed computation [22] that in turn were inspired by works in distributed decision making examined by statisticians and economists [13], [198],[213].

Space-borne optical interferometry is an active area of research for a number of future scientific missions by NASA, such as the Terrestrial Planet Finder [3] and by the European Space Agency, such as the Darwin Mission [1]. Interferometry is one of the cornerstones of applied optics [32]; for the spaceborne application of interferometry, see [224]. Molecular switch and tautometers are of great interest in nanotechnology, examples of which can be found in [146],[172],[206]. Social networks is an active area of research in sociology, statistics, and economics; see for example, Wasserman and Faust [241]; for a more network-centric treatment, see the books by Goyal [105] and Jackson [122].

For complementary references related to this book, with somewhat different emphasis and outlook, see the books by Ren and Beard [204], and Bullo, Cortés, and Martínez [41].