



A Critical Review of Communications in Multi-robot Systems

Jennifer Gielis¹ · Ajay Shankar¹ · Amanda Prorok¹

Accepted: 18 July 2022 / Published online: 22 August 2022
© The Author(s) 2022

Abstract

Purpose of Review This review summarizes the broad roles that communication formats and technologies have played in enabling multi-robot systems. We approach this field from two perspectives: of robotic applications that *need* communication capabilities in order to accomplish tasks, and of networking technologies that have enabled newer and more advanced multi-robot systems.

Recent Findings Through this review, we identify a dearth of work that holistically tackles the problem of co-design and co-optimization of robots and the networks they employ. We also highlight the role that data-driven and machine learning approaches play in evolving communication pipelines for multi-robot systems. In particular, we refer to recent work that diverges from hand-designed communication patterns, and also discuss the “sim-to-real” gap in this context.

Summary We present a critical view of the way robotic algorithms and their networking systems have evolved, and make the case for a more synergistic approach. Finally, we also identify four broad *Open Problems* for research and development, while offering a data-driven perspective for solving some of them.

Keywords Communication · Multi-robot systems · Robot networks

Introduction

The use of multiple, *connected* robots in the place of individually uncommunicative robots provides evident gains by facilitating the inter-robot coordination that allows for work distribution, spatial coverage, and specialization. An increasing variety of applications leverage such networks of robots, including logistics [1, 2], resource distribution [3, 4], transport systems [5–7], manufacturing [8], and agriculture [9, 10].

These applications depend on an orchestration of robots over time and space that allows them to jointly

work towards common higher-order goals, to deconflict individual actions in shared environments, and to share information in distributed computing schemes. Communication and the mutual exchange of information (state and control) are key to facilitating such interactions.

Early work in the multi-robot domain drew from nature-inspired paradigms [11], and consequently focused on devising collective behaviors that depended purely on *local* interactions of robots in close proximity [12]. A variety of transmission media (e.g., infrared) are used for such near-field communication schemes [13, 14]. Other nature-inspired work built on *implicit* communication and self-organization through stigmergy, by which robots coordinate indirectly through traces left in the environment [15]. The benefits of such peer-to-peer *decentralized* communication paradigms are manyfold, in particular due to their inherent robustness and scalability [16]. *Centralized* radio-based communication architectures have become increasingly popular in various instances, especially when the task requires performance guarantees; representative applications include product pickup and delivery [17], item retrieval in warehouses [3], and mobility-on-demand services [18]. Improvements in communication technologies, both hardware and software, have furthered more data-intensive applications, such as cloud robotics [19, 20].

Jennifer Gielis and Ajay Shankar contributed equally to this work.

This article is part of the Topical Collection on *Group Robotics*

✉ Jennifer Gielis
jag233@cl.cam.ac.uk

Ajay Shankar
as3233@cam.ac.uk

Amanda Prorok
asp45@cam.ac.uk

¹ Department of Computer Science and Technology,
University of Cambridge, Cambridge, UK

Explicit communication methods generally assume that robots can broadcast information within a local neighborhood that comprises of tens to hundreds of individuals, or that a fixed network infrastructure is available. Yet, in reality, densely populated workspaces adversely affect communication capabilities because of practical contention over channel bandwidth and airtime [21]. Such networks are additionally burdened by clutter that can induce signal fading, leading to a drastic decrease in the expected communication range. This problem is compounded by the need for real-time transmission requirements in highly dynamic robot networks. Indeed, topologies and capabilities *demand*ed by robotic applications are practically hostile to radio performance (because these radio networks were not initially designed with robotics in mind). As a consequence, the vast majority of robot applications are designed to merely work around available network technologies, and optimize their performance within the given constraints.

Our review is motivated by a lack of studies that provide a high-level overview of the interplay between communication networks and their role in robotic applications. Figure 1 graphically demonstrates a typical architecture of a multi-robot control scheme, in which each robotic system is designed strongly around its perception and control strategies. As mentioned prior, robot control algorithms generally do not actively employ the output of the communications network as part of the control loop, and as a result often overlook factors such as network contention. The result is a wide array of optimizations that work in favor of the network, but often not for autonomy, or vice-versa. Hence, we argue that a better co-optimization scheme (illustrated on the left in Fig. 1) would consider all aspects of the architecture simultaneously. The illustration shows this scheme

linked via an “oracle”, which are sources of error estimations from incomplete information, that facilitates this, i.e., given a hypothetical oracle, we posit that one can co-design the various algorithm layers on the right.

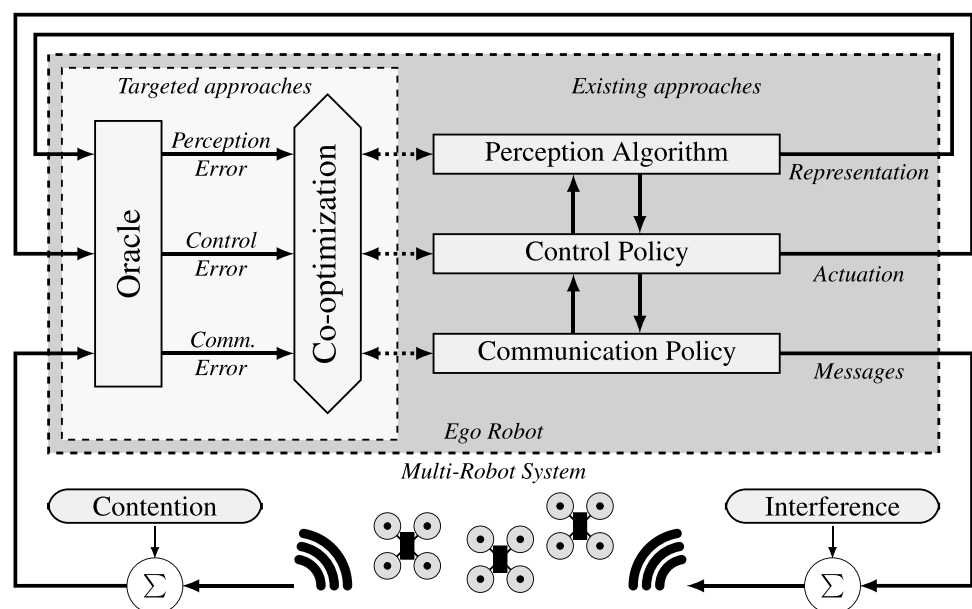
In this survey, we capture a variety of network architectures and technologies, and a variety of multi-robot applications that employ them. A careful choice of communications architecture, medium and algorithm is key to ensuring that a given robot task can be completed. Therefore, we will also explore some of the newer approaches that consider bypassing such hand-crafted selections, and attempt to model inter-robot communications in a data-driven fashion.

Factors Influencing Robot Network Design

Choosing an adequate communications architecture, medium, and algorithm is key to ensuring desired robot performance. In the following, we distill the factors that influence the robot network design choices. We elaborate upon them in the following four categories: (i) application, (ii) robot, (iii) algorithm, and (iv) environment, and give an illustrative example for each.

The Application: The application defines *what the shared information is for*, and *how the robots need to interact* to solve the problem at hand. *Examples:* Real-world applications such as in environmental monitoring and agriculture require groups of robots to act over large distances (often operating with robots separated by $\sim 1000\times$ body lengths). Such sparsely distributed robot systems, hence, necessitate networking capabilities that can span larger spaces [22•]. Other applications, such as cooperative driving [5], formation control [23], and flocking [24]

Fig. 1 A flow diagram showing how perception, control and communications typically interact in a multi-robot system. On the left is a proposed future approach to co-optimization, discussed in Section [Challenges and Open Problems](#), which now admits communications and its confounding factors as integral to any co-optimization strategy designed for a multi-robot system



require uninterrupted, situated, close-range communication for tight inter-robot coordination and control.

The Robot: The robot (and the physical hardware) define local *constraints on the frequency and format* of information to be transmitted and received. *Examples:* A quadrotor that uses state information for local stability control requires an update frequency in the order of several hundred Hertz; while on-board IMUs can provide the necessary information for body stabilization, extrinsic pose estimates are still required for tasks, and must be received at relatively high rates (e.g., 100 Hz) [23, 25]. Lack of reliable updates naturally poses a significant risk to tasks that require tight coordination, such as outdoor flocking and formation control; while sparse outdoor flight has been demonstrated in a team of 30 drones [26], there is a dearth of results on *dense and agile* outdoor flight. Moreover, in GPS-denied environments, robots resort to on-board sensing, and consequently, require dependable inter-robot communication to achieve group behavior.

The Algorithm: The algorithm connects the *application* to the *robot*, and essentially sets conditions on the nature of information that needs to be received (e.g., global or local), and when (e.g., asynchronously or synchronously, and how often). *Examples:* In allocation problems, the optimization objective is often global, and to achieve optimality, we deploy centralized algorithms that collect all robot-to-task assignment costs (e.g., expected travel times) to determine the optimal assignment (e.g., by running the Hungarian algorithm) [27, 28]. Similarly, multi-robot path planning has an optimal solution (for both makespan and flowtime objectives), but only when the computational unit has access to full system information [29]. In the absence of full observability, robots need to resort to locally available knowledge. In *decoupled* prioritized path planning, robots communicate to mutually deconflict their path plans in time-space [30–32]. Each time a robot’s plan changes, its robot neighborhood changes, or a new conflict arises, the deconfliction process restarts.

The Environment: The environment defines *under what conditions* shared information is delivered. *Examples:* Are the robots operating indoors or outdoors, or both [33]? Does the workspace afford a fixed (and possibly centralized) communications infrastructure, or must we instead rely on ad hoc networking? Is the environment cluttered with obstacles that interfere with wireless signals? What medium can we use, e.g., are the robots operating in air, under water, or in space? What legal jurisdictions regulate the communication infrastructure? And finally, is the communication channel safe, or can it be spoofed [34], or robots attacked [35, 36]?

Communication Schemes

In this section we discuss multi-robot communication from the perspective of the underlying communications technologies, focusing upon the challenges, limitations and optimizations that are relevant in multi-robot system networks. Figure 2 shows a timeline with key wireless communication mechanisms, and some representative multi-robot applications that they enabled for the first time.

Challenges

Synchronicity. Specifying robotic data flows is often the first consideration in discussing the challenges of a wireless data protocol. For example, it is often implicitly assumed that multi-robot control algorithms are executed synchronously by every participant [47]. This introduces a hard timing constraint on the maximum allowable delay in message delivery between those participants. This is, however, a feature that commonly deployed communications protocols are not designed to meet, with “best-effort” message delivery being the standard paradigm [48].

Dynamic Topologies. Hard timing constraints are often exacerbated by highly connected communications topologies that are dynamic, where a robot must communicate its status with many different (or sometimes every) participating robot(s). This can lead to a high degree of contention for radio resources since many messages may need to be sent at every control loop. While there are schemes that aim to minimize redundant data transmissions (see Section [Communication-Aware Algorithms](#)), it remains true that, as multi-robot networks increase in scale, communications technologies must be selected and designed specifically to manage the dynamics of the application [49], something that is generally overlooked in robotic networks today.

Message Frequency. Bandwidth is often employed as a metric to specify the demands on a communications link [50]. However, this is often an insufficient characterization by itself, since the underlying technology may have significant overheads per message, and robot teams often depend on low-latency messaging as well. This is particularly true for ad hoc networks where there is no central entity enforcing message scheduling, to the extent that many communications protocols will not approach their rated bandwidths in highly connected ad hoc topologies, where overheads such as contention dominate radio resource consumption [21].

Connectivity. Since a greater connectivity range implies reachability and information exchange with more robots, it has an obvious impact on the overall messaging rate any specific robot must handle. The spatial density of

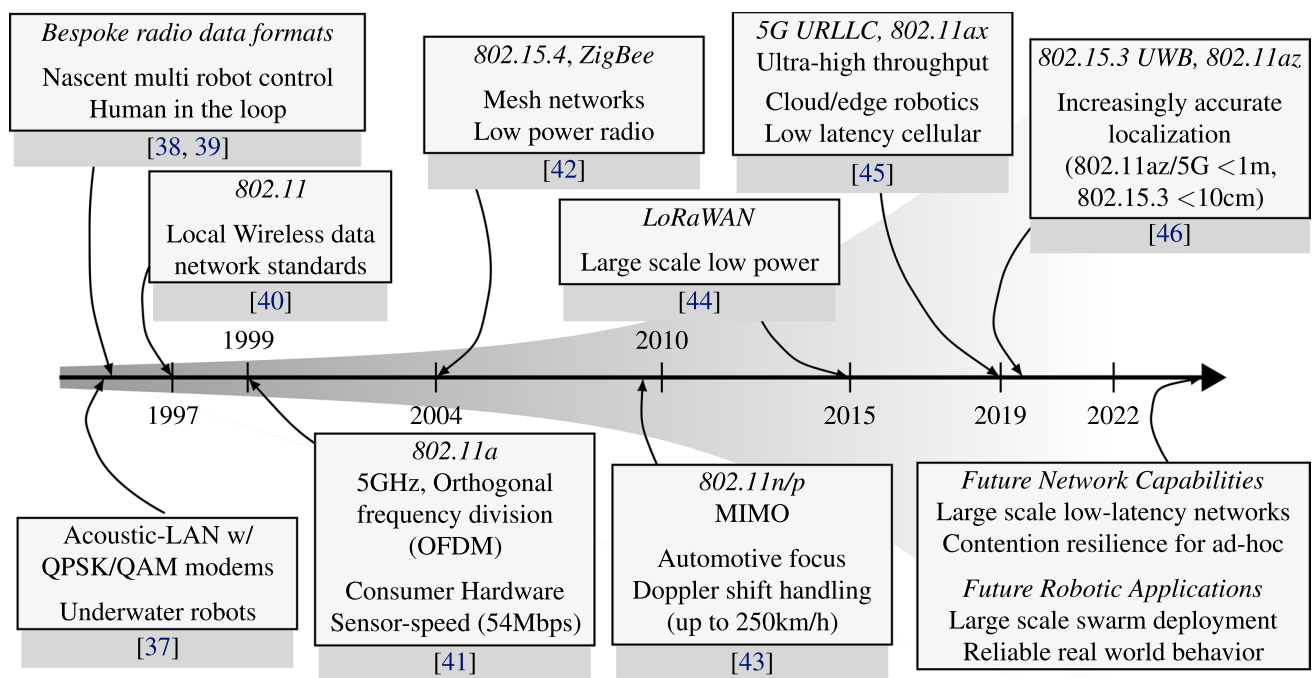


Fig. 2 An abridged timeline showing some key wireless communication mechanisms for robots. The shaded area represents the magnitude of theoretical capabilities (bandwidth or inverse latency), which

have been increasing super-linearly since 2010. Boxes show communications standards or technologies, and their specific properties that are useful in the multi-robot control space

robots must be considered while discussing range, and two key factors of interest emerge as a consequence. Firstly, as ranges increase, radio-based links are more prone to fading and interference [51] even when transmission power is commensurately increased. Secondly, robot control algorithms often assume a fixed range [52], which may result in greater inter-connectivity, and sometimes increased messaging rates in dense scenarios.

Dynamic Routing. The case where the required range of communication exceeds the underlying capabilities of the radio hardware onboard must also be considered, as this implies a mesh-type network where a message must traverse multiple robots (network nodes). This first requires planning the robots' paths (discussed in Section [Communication-Aware Planning](#)), before accounting for the computational and protocol overheads of robots processing messages other than their own. Then, the problem is that of message routing decisions and dynamic topologies. The routing decision problem is generally central to ad hoc mesh networks; this is only made more challenging by the potential for rapid shifts in communications topology, especially in highly mobile, or large-scale robotic scenarios [53]. Hard timing guarantees, in the range required by robotic control, are not currently available at non-trivial scales (especially with multi-hop routing over dynamic topologies), though some attempts have been made in this direction [54].

Operational Environment. Robotic networks will invariably be required to operate in environments with external noise and interference, which cause unpredictable impacts on link quality. This informs the selection of communications protocols, since some protocols operate in a licensed spectrum with reduced external interference, or are otherwise less prone to external noise due to atmospheric attenuation (60GHz). Doppler shift requires similar consideration, because many communications technologies fail at high relative velocities. Generally speaking, protocols that depend upon fine-grained frequency division multiplexing are more prone to Doppler related errors [55], and such schemes are often used in high bandwidth techniques.

Communications Scheme Selection

Despite considerable research interest, there are no current wireless data standards explicitly designed for exchanging information between autonomous robots [56]. Currently deployed robot-to-robot networks (such as [26]) depend upon more generic wireless data networking standards which are not typically optimized for the challenges discussed above. In the absence of a specific standard, we will discuss the strengths and weaknesses of existing technologies for the multi-robot control application.

Ad hoc networks map well to the communications patterns required of decentralized robotic control, and the most relevant for this survey are Mobile Ad Hoc Networks (MANETs) which deal with the problem of facilitating communication between mobile nodes without coordination from infrastructure [57]. More specific forms of interest are Vehicular Ad Hoc Networks (VANETs) [58] and Flying Ad Hoc Networks (FANETS) [59], where the former generally deal with automotive use cases and the latter drones and UAVs, and these are more exposed to dynamic conditions that are expected from robotic ad hoc networks. Local area networking technologies are well suited for ad hoc networking.

In contrast with ad hoc operation, infrastructure networks map more closely to a centralized robotic control, where communications patterns are more similar to traditional bandwidth-focused networking applications. Despite this, hard latency requirements and rapid robot movement require specific mechanisms at the protocol level, which are not common for either cellular or local area standards.

Local Area Networks

The IEEE 802.11 protocol suite, commonly known as “wi-fi,” is frequently used due to abundant hardware availability, IP networking interoperability, high data rates and license-free operation. It is also capable of both infrastructure and ad hoc operations, which simplifies deployment from laboratory environments into the real world. The failures of 802.11 become apparent during such deployment processes [60, 61], because larger ranges, robot counts and velocity-induced Doppler shift cause lower message delivery rates than are seen in static 802.11 deployments. 802.11p, and its successor 802.11bd, have both introduced specific modifications to the *physical* layer [43] to more robustly handle both range and Doppler-induced problems for VANET use cases, which potentially transfer to robotic control as well.

IEEE 802.15.4 has been commonly used as the basis for a number of different higher-layer protocols, including ZigBee. It has been deployed in the context of wireless sensor networks and multi-robot systems [42] due to hardware availability, license-free operation, low power usage, and flexible communications models that permit both IP-based and more simplified serial-like messaging. The major drawbacks are relatively low range and data rates. LoRaWAN [44] is an attractive alternative that maintains the positive aspects of 802.15.4 for the robotic use cases, but with a focus upon long range transmission (up to 16km) and a physical layer that is resilient to Doppler errors; however, its communications model is infrastructure based.

Both 802.11’s and 802.15’s underlying dependence upon the CSMA/CA collision avoidance scheme allows for the minimization of contention related losses without

an authoritative central scheduler [62]; however, they have unbounded maximum latency on message delivery, and reduced message delivery rates with higher numbers of robots on the network. LoRaWAN uses a pure ALOHA protocol mechanism [63], and therefore scales even more poorly than the IEEE schemes. These characteristics make these protocols unsuitable for deployment on robots without modifications; fortunately, there are techniques proposed in the literature to help make these technologies more scalable [64, 65].

All of the protocols mentioned within this section share similar routing problems when it comes to highly dynamic topologies, in that they depend upon a network-layer routing mechanism to direct traffic without reliably converged information about the current disposition of other nodes. This issue is well covered by [66], which categorizes and surveys many of the different approaches towards this routing problem. [67] includes the routing issue amongst a general survey of the issues in UAV networking.

Cellular Networks

Cellular networks avoid many of the problems encountered by local networking standards by making use of a nearly universal infrastructure-based communications model, licensed radio spectrum access, as well as economics of scale, all of which make extremely complex base station hardware and protocols commonplace. The centralized message scheduler and sophisticated radio resource management [68, 69] are significantly more scalable than typical ad hoc networks. These characteristics appear to be a good fit for robot control, however, the financial cost of network access, coupled with limited flexibility in logical network configuration have limited robotic deployments outside controlled environments.

For decentralized robot control, peer to peer traffic is routed through the infrastructure, inducing a minimum latency overhead [70] that could exceed timing constraints. 4G in particular has an access latency on the order of 50ms [60]. Additionally, cellular standards are naturally dependent upon the presence of infrastructure, which cannot always be assumed. Furthermore, the radiation pattern of cellular networks is typically setup assuming ground based users, and so aerial robots could experience degraded performance due to leaving the vertical coverage of cell antennas [71].

4G LTE supports direct Device-to-Device (D2D) modes that permit devices to communicate with each other in a local region by reserving some subset of radio resources in their local area from the network operator. LTE-V2V is a variant of this specifically for automotive use cases [72]. This avoids the overhead of using the infrastructure as a relay but also has a cost in minimum association time and is dependent upon the network operator ceding resources on

demand. Some proposals extend D2D cellular radio techniques into the unlicensed spectrum, or specifically, licensed sub-bands [73], however, this also has not been widely used in real-world systems due to the recency of the specification and a dearth of capable hardware.

5G introduced the Ultra-Reliable Low Latency Communications (URLLC) service to address the issues with low-latency medium access [45] in a centralized or decentralized manner, with a variety of proposed physical layer [74] and mac layer [75] techniques. Despite the promise of these approaches, as well as 3GPP release 15 (including URLLC) being released in late 2017, roll out of these technologies into real-world networks has been limited, and considerable research is ongoing as to the best implementation methods for 5G's technical goals [76].

Hybrid Schemes

Due to advancements in radio hardware, the most recent 802.11 revision, 802.11ax, has specified a physical radio resource allocation scheme that is far more similar to cellular standards, with multiple fine-grained frequency divisions being available within a single logical channel. Many proposals have been made to have future cellular standards directly inter-operate with cellular networks to leverage the best capabilities of both [77]. For robot networks, this may prove to be highly valuable, permitting human control overrides over cellular infrastructure and low-latency robot-to-robot communications with a unified logical network addressing system for easier lab-to-world deployment, and efficient operation in control schemes that have evolving requirements throughout a single deployment. Though exciting, these proposals are still in their nascent phase and many issues remain to be addressed.

Though the state-of-the-art use of OFDMA in both 5G and 802.11ax significantly alleviates the contention problem due to the larger number of transmission slots made available, extremely dense ad hoc robot networks may still run into the limit of CSMA/CA. Non-orthogonal media access (NOMA) is a very promising technology that has the possibility to further extend the effective simultaneous radio resources available [78], and therefore reducing the contention problem. In cellular systems, the network operator still has authoritative control over their radio resources, and so, grant-based schemes induce significant overhead despite the reduced contention; though grant-free access has attracted significant attention [79]. Even in grant-free schemes, NOMA still requires coordination to ensure that node configurations do not overlap, and hence, message loss due to contention remains an open problem in decentralized networks without a coordinating infrastructure.

Communication-Aware Algorithms

Regardless of which underlying communication scheme or protocol is employed, unlimited and unconstrained communication cannot be assumed for any interactive scenario. A significant amount of literature in multi-robot applications, however, has generally focused on designing control schemes that do not explicitly model this dependency. This is reflected in the vast majority of literature in robot flocking [80–83]. We argue that the problem becomes more pronounced in cases where the robots need to deconflict and replan their motions in tight and constrained spaces [81, 82]. While some consideration for communication asynchronicity is made in some of the more recent works [82], the challenge is generally far from being solved.

One straight-forward approach to handle this is to simply *reduce* the amount of data (frequency, packet-size etc.) that needs to be communicated between agents. In an exploration problem, this is often done through various novelty metrics that determine whether a new data-point needs to be communicated [84]. Trawny et al. [85] have proposed localization estimators that perform well by quantizing the transmitted information to very small packets, thereby tackling severe link constraints.

On the other hand, there is also a sustained research interest in modeling the communication channels between the agents, and factoring that as a constraint into the motion planning problem. This is done primarily to ensure robustness of a control scheme against imperfect and noisy communications. Alternatively, planning schemes have also considered communication as a sub-task (almost as if “scheduling” communications at intervals). Finally, there are several approaches that consider a joint optimization scheme, where path planning and communication planning are carried out in tandem. We divide this body of work into these three broad styles.

Communication-Aware Planning

As mentioned earlier, planning robot motions or trajectories that consider some model of the underlying communication links is an active area of research. Mularidharan and Mostofi provide a comprehensive overview of such methods [86]. For instance, several authors have considered the task of coverage & formation control by a team of robots. Evidently, these domains require explicit factoring of communication constraints into the planning problem [84, 87]. One way this is approached is by analyzing the stability of a formation under various communication link latencies [88]. This can then be then integrated into the control problem for a more reliable system [88, 89] that

is “aware” of such latencies. Formation control laws are also explored that allow agents to maintain some degree of coordination while respecting limited communication ranges of their neighbors [90, 91]. This approach is also sometimes utilized in the context of cooperative target localization [92], or under constraints of 3G/4G mobile networks [93].

Path planning has also been developed such that connectivity with a subset of base-stations [94], or with some agents [95, 96] is maintained. Similar methods that plan for multiple robots (with connectivity constraints) involve using ACO-based (ant colony) planning [97] or genetic algorithms [98].

Plan-Aware Communications

In heavily constrained spaces, it is often desirable to design a network architecture that considers the planned path, and seeks opportunities to communicate therein. Underwater robots, for instance, have very limited communication capacities [99, 100], and this is an active field of interest; Zolich et al. [101] provide a comprehensive survey on the various challenges and solutions. Hollinger et al. have considered scheduling algorithms for underwater robotic sensor networks and show how path planning algorithms depend on these [102, 103]. Since bandwidth and interference constraints are much more severe in these environments, such scheduling algorithms often model the *value* of communicating at a particular timestep [103, 104]. This also plays a role in determining whether communicating has a positive impact on the state of the robot system [105, 106], and is also studied as an online decision problem [107], and an optimization problem that considers when/what to communicate [108]. Recent developments in subterranean robots operating under severe communication constraints have also explored the strategy of developing/maintaining communication “backbones” for explorer robots to continue on frontiers [109, 110], and of explicitly splitting communication pathways by plan and priority [111].

Joint Planning

Several of the works listed in the previous subsections may also be seen as jointly optimizing for communication quality as well as path qualities. However, there are other approaches that attempt to explicitly model this optimization problem. For instance, Kantaros and Zavlanos [112] propose a scheme that *alternates* between the two optimization problems sequentially. The nature of this scheme often makes it difficult to prove hard guarantees regarding optimality; however, a more hybrid approach in which the two controllers interact can offer more guarantees on network integrity available data rates [113]. A joint optimization scheme, on the other hand, can formulate this problem well; for instance,

using an LQ (linear-quadratic) form can additionally offer robustness guarantees as well [114]. Yet another means of joint optimization is to consider the system as a cyber-physical system (CPS), where the “cyber” controller handles the communications domain, and the “physical” controller handles the kinematics of the robot [115]. Such models allow designers to factor various other elements of a CPS system, such as dynamically adapting one of the subsystems (communication capacity) while still maintaining the coupling with the other [116, 117].

Leveraging Machine Learning for Communication

Designing bespoke, hand-crafted communication protocols and behaviors is tedious and difficult. Firstly, numerous works point to the hardness of synthesizing decentralized policies (that have to operate in a partially observable regime), even when a centralized template is known [118, 119], and they leave the question of *how* (what, when, and to whom) to communicate unanswered. Secondly, the vast majority of existing robot communication strategies are based on idealistic operational assumptions, and besides a few specialized approaches to dealing with message loss, delay, or corruption, e.g., [34, 35, 120], it is not at all clear how to approach such problems in a manner that is transferable across applications. Leveraging machine learning methods is a promising new avenue to tackle some of these challenges.

Learning Communication Mechanisms

Message routing decisions in robotic mesh networks are complicated by highly dynamic topologies. While many routing mechanisms exist in ad hoc networks, these generally depend upon relatively slowly changing network conditions to function effectively. Many manually specified heuristic methods exist [66•]; however, these may lead to sub-optimal decisions as they may be constructed upon incorrect assumptions about the target network environment. Learning methods provide an attractive alternative, and have been explored in some depth in routing generally [121]. An interesting example in the context of FANETs can be found in Zheng et al. [122], who propose RLSRP which applies an online reinforcement learning method to the routing decision problem and shows improved performance across several metrics, including delivery latency.

Channel modeling and resource allocation are also key networking problems that are challenging for first principles methods to solve that can be improved with learning [123]. Unsupervised learning has been applied to channel modeling, which allows for the optimization of transmission

power by accurately estimating the quality of links to other network participants [124].

Learning Communication Behaviors

Learning-based methods have proven effective at designing robot control policies for an increasing number of tasks [125, 126]. Recent work utilizes a data-driven approach to solve multi-robot problems, for example for multi-robot motion planning in the continuous domain [127] or path finding in the discrete domain [128].

Yet, research on learning how to *synthesize robot-to-robot communication policies* is nascent. From the point of view of an individual robot, its local decision-making system is incomplete, since other agents' unobservable states affect future values. While the manner in which information is shared is crucial to the system's performance, the problem is not well addressed by hand-crafted (bespoke) approaches. Learning-based methods, instead, promise to find solutions that balance optimality and real-world efficiency, by bridging the gap between the qualities of full-information centralized approaches and partial-information decentralized approaches [129].

Key to the decentralization of centralized (optimal) policies is the property of *permutation equivariance*. Permutation equivariance ensures that at the robot network level, the set of actions automatically rearranges itself as the agents swap order. One of the earliest works that satisfy this property is [130]. This was concurrently developed by a line of work that builds on Graph Neural Networks (GNNs), which are permutation equivariant by design [131–133]. GNNs have since then shown promising results in learning explicit communication strategies that enable complex multi-agent coordination [134–137].

When deploying GNNs in the context of multi-robot systems, individual robots are modeled as nodes, the communication links between them as edges, and the internal state of each robot as graph signals. By sending messages over the communication links, each robot in the graph indirectly receives access to the global state. A key attribute of GNNs is that they compress data as it flows through the communication graph. In effect, this compresses the global state, affording agents access to relevant encodings of global data. Since encodings are performed locally (with parameters that can be shared across the entire graph), the policies are intrinsically decentralized. In cases where the downstream task is tightly coupled with the communication requirements, it is beneficial to optimize the communication strategy jointly with perception and action policies. This was done in [138], for multi-robot flocking, and in [137], for multi-agent path planning. These frameworks implement a cascade of a convolutional neural network (CNN) and a GNN, which they jointly train so that image features and

communication messages are learned in conjunction to better address the specific task. Recent work also shows how GNNs can be augmented by *attention modules* to produce *message-aware* communication strategies that allow robots to discern between important and less important message elements [139].

Approaches from within the multi-agent reinforcement learning (MARL) community tackle the learning of continuous communication protocols by formulating the problem as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) [140–142]. The work in [143•] learns a targeted multi-agent communication strategy by exploiting a signature-based soft attention mechanism (whereby *message relevance* is learned). Similarly, the work in [144] has each robot learn to reason about other robots' states and to more efficiently communicate trajectory information (i.e., when and to whom), and applies the solution to the problem of collision avoidance. While efficient *cooperative* communication strategies are desirable, the work in [145] shows how separate robot teams can learn to communicate with *adversarial* behaviors that contribute to manipulative (non-cooperative) behaviors. Clearly, underlying training paradigms need to be carefully designed to avoid such outcomes.

Challenges and Open Problems

We finally present some avenues of research and engineering that are worth exploring in order to address our critiques discussed so far. We categorize them into four broad *Open Problems*.

1. Co-design. An emergent theme throughout this survey is the lack of approaches that co-design the robot and its communication capabilities. A variant of this concept [146, 147] considers a basic parallel reconfiguration of a network as well as the robot's controller that can be beneficial when the robot moves across network stations. However, a true co-design scheme will jointly evolve all layers of the networking stack to favor the robotic task at hand. Design of a meta-system that is able to compute the limitations of robotic requirements as well as network capabilities and dynamically throttle both may be essential to safe deployment of robots into the real world. Any robotic control algorithm that uses explicit communications is vulnerable to failure if the network unexpectedly under-delivers, and performs sub-optimally if the network over-delivers — managing this resource allocation problem in a real-world multi-robot setting is a subject we will tackle in our future work.

2. Data-driven optimization. Machine learning, and specifically, reinforcement learning, can drive the development of multi-robot communications into new and interesting paradigms. Existing approaches that already learn what/when to send (and whom to send to) [130, 139, 144] still

often depend on hand-designed architectures and specific task groups. With sufficiently large datasets, novel machine learning architectures also have the potential to learn to optimize multiple aspects of multi-robot systems at once (e.g., perception, action and communication [138]).

3. Sim-to-real for robot networks. The problems in sim-to-real transfer of robot coordination strategies are generally exacerbated by the “reality gap” found in communications [129]. Practical communication links suffer from message dropouts, asynchronous and out-of-order reception, and decentralized mesh topologies that may not offer reliability guarantees. Since multi-robot policies are typically trained in a synchronous fashion, these factors are hard to capture and simulate [148]. Furthermore, very few studies have captured any of these network effects in a large-scale setting [21]. Consequently, we find that embedding the reality gap of robot networking into data-driven approaches to multi-robot planning is an open research domain.

4. New technologies/schemes. As discussed in Section Communications Scheme Selection, there is a need for wireless data standards that specifically target the communication requirements of connected robots. The IEEE 1920 working group is a significant step in this direction, which was formed to propose a protocol that is intended for autonomous robotic networks [56]. Such a protocol is likely to be founded on 802.11bd since it is already a significant leap forward [43] from the legacy 802.11p standard used in V2V standards today.

Additionally, future 5G updates and 6G cellular communications promise dramatic improvements that hold the potential to bring cloud- and edge-computing at the forefront of many data-intensive multi-robot collaborations. Finally, we also note that geographic routing in FANETs may be an enabling technology for practically dealing with highly dynamic routing topologies. This will, however, require holistic developments in robot control algorithms that work in tandem to avoid an additional information distribution problem.

Conclusion

Through this manuscript, we have presented a survey of communication technologies and their role in enabling multi-robot applications. We have broadly covered the various technologies that have played key roles in networked robotics, and have also discussed how state-of-the-art robot applications typically deal with network constraints. Our approach to this has been mostly critical, and thus, has identified several deficiencies in the way robotics and networks have evolved. Towards the end, we also cover machine learning approaches and their role

in developing data-driven communication strategies. We conclude the article with a list of challenges and open problems that the community currently faces, and also provide an outlook for how learning-based approaches can tackle several of them.

Funding Jennifer Gielis was supported by an EPSRC Doctoral Training studentship. The authors also received support from ARL DCIST CRA W911NF-17-2-0181 and the European Research Council (ERC) Project 949940 (gAIA).

Declarations

Conflict of Interest The authors declare no competing interests.

Human and Animal Rights and Informed Consent This article does not contain any studies with human or animal subjects performed by any of the authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Papers of particular interest, published recently, have been highlighted as:

- Of importance

1. Tilley J Automation, robotics, and the factory of the future. McKinsey Market Report; 2017
2. Kamagaew A, et al. Concept of Cellular Transport Systems in facility logistics. In: Proceedings of the 5th international conference on automation, robotics and applications. IEEE; 2011 <https://doi.org/10.1109/ICARA.2011.6144853>
3. Enright JJ, Wurman PR. Optimization and coordinated autonomy in mobile fulfillment systems. Workshops at the Twenty-Fifth AAAI conference on artificial intelligence; 2011
4. Ma H, et al. Lifelong multi-agent path finding for online pickup and delivery tasks. 2017. [arXiv:1705.10868](https://arxiv.org/abs/1705.10868) [cs]
5. Hyldmar N, He Y, Prorok A. A fleet of miniature cars for experiments in cooperative driving. IEEE International Conference Robotics and Automation (ICRA); 2019. <https://doi.org/10.17863/CAM.37116>
6. Dressler F, Hartenstein H, Altintas O, Tonguz O. Inter-vehicle communication: Quo vadis. In: IEEE Communications Magazine; 52.6. 2014. pp. 170–77
7. Ferreira M, et al. Self-organized traffic control. In: Proceedings of the seventh ACM international workshop on VehiculAr Inter-NETworking. ACM; 2010. pp. 85–90

8. Cherubini A, et al. Collaborative manufacturing with physical human-robot interaction. *Robotics and Computer-Integrated Manufacturing*. 2016;40:1–13.
9. Noguchi N, Barawid OC. Robot farming system using multiple robot tractors in Japan agriculture. In: *IFAC Proceedings volumes*; 2011. 44.1. pp. 633–37. <https://doi.org/10.3182/20110828-6-IT-1002.03838>
10. Albani D, et al. Monitoring and mapping with robot swarms for agricultural applications. In: *14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. 2017
11. Bonabeau E, et al. *Swarm intelligence: from natural to artificial systems*. Oxford University Press; 1999
12. Nagpal R, Shrobe H, Bachrach J. Organizing a global coordinate system from local information on an ad hoc sensor network. In: *Information processing in sensor networks*. Springer; 2003. pp. 333–348
13. Rubenstein M, Cornejo A, Nagpal R. Programmable self-assembly in a thousand-robot swarm. In: *Science*. vol 345.6198; 2014. pp. 795–99
14. Pugh J, Martinoli A. Relative localization and communication module for small-scale multi-robot systems. In: *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006. IEEE*; 2006. pp. 188–93
15. Beckers R, Holland OE, Deneubourg J-L. From local actions to global tasks: stigmergy and collective robotics. In: *Prerational Intelligence: Adaptive Behavior and Intelligent Systems Without Symbols and Logic, Volume 1, Volume 2 Prerational Intelligence: Interdisciplinary Perspectives on the Behavior of Natural and Artificial Systems, Volume 3*. Springer; 2000. 1008–22
16. Brambilla M, et al. Swarm robotics: a review from the swarm engineering perspective. In: *Swarm Intelligence*. Springer; 2013. vol 7.1. pp. 1–41
17. Grippa P, et al. Drone delivery systems: job assignment and dimensioning. In: *Autonomous Robots*; 2019. vol 43.2. pp. 261–74. <https://doi.org/10.1007/s10514-018-9768-8>
18. Spiesser K, et al. Shared-vehicle mobility-on-demand systems: a fleet operator’s guide to rebalancing empty vehicles. In: *Transportation Research Board*. 2016
19. Chinchali S, et al. Network offloading policies for cloud robotics: a learning-based approach. In: *Autonomous Robots*. Springer; 2021. vol 45.7. pp. 997–1012
20. Kehoe B, et al. A survey of research on cloud robotics and automation. In: *IEEE Transactions on automation science and engineering*. IEEE; 2015. vol. 12.2. 398–409
21. Gielis J, Prorok A. Improving 802.11 p for delivery of safety-critical navigation information in robot-to-robot communication networks. In: *IEEE communications magazine*. IEEE; 2021. vol 59.1. pp. 16–21
22. Tarapore D, Groß R, Zauner K-P. Sparse robot swarms: Moving swarms to real-world applications. In: *Frontiers in Robotics and AI*. vol 7. 2020. pp. 83. **This paper outlines some of the key challenges in communications and control when robotic teams are spread over larger volumes (as opposed to dense operations with potentially lower latencies.)**
23. Preiss JA, et al. Crazyswarm: A large nano-quadcopter swarm. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE; 2017. pp. 3299–304
24. Tolstaya E, et al. Learning decentralized controllers for robot swarms with graph neural networks. In: *Conf. Robot Learning 2019*. Osaka, Japan. *Int. Found. Robotics Res*; 2019
25. Kushleyev A, et al. Towards a swarm of agile micro quadrotors. In: *Autonomous Robots*; 2013. vol. 35.4. pp. 287–300
26. Vásárhelyi G, et al. Optimized flocking of autonomous drones in confined environments. In: *Science Robotics*; 2018. vol. 3.20. p. eaat3536
27. Prorok A Redundant Robot Assignment on Graphs with Uncertain Edge Costs. In: *Distributed Autonomous Robotic Systems*. Ed. by Correll, N., Schwager, M., and Otte, M. Springer Proceedings in Advanced Robotics; 2019, pp. 313–27. ISBN: 978-3-030-05816-6.
28. Khamis A, Hussein A, Elmogy A. Multi-robot task allocation: A review of the state-of-the-art. In: *Cooperative robots and sensor networks 2015*. 2015. pp. 31–51
29. Yu J, LaValle SM. Optimal multi-robot path planning on graphs: Complete algorithms and effective heuristics. 2015. [arXiv:1507.03290](https://arxiv.org/abs/1507.03290) [cs]
30. Wu W, Bhattacharya S, Prorok A. Multi-robot path deconfliction through prioritization by path prospects. In: *IEEE International Conference on Robotics and Automation*; 2020. [arXiv:1908.02361](https://arxiv.org/abs/1908.02361)
31. Cáp M, et al. Asynchronous decentralized prioritized planning for coordination in multi-robot system. In: *Intelligent robots and systems (IROS), 2013 IEEE/RSJ international conference on*; 2013. pp. 3822–29
32. Desaraju VR, How JP. Decentralized path planning for multi-agent teams with complex constraints. In: *Autonomous Robots*; 2012. vol. 32.4. pp. 385–403
33. Dong J, Distributed real-time cooperative localization and mapping using an uncertainty-aware expectation maximization approach. In: *et al. IEEE international conference on robotics and automation (ICRA)*. IEEE. 2015;2015:5807–14.
34. Gil S, et al. Guaranteeing spoof-resilient multi-robot networks. In: *Autonomous Robots*; 2017. vol. 41.6. pp. 1383–1400. <https://doi.org/10.1007/s10514-017-9621-5>
35. Saulnier K, et al. Resilient flocking for mobile robot teams. In: *IEEE Robotics and Automation Letters*; 2017. vol. 2.2. pp. 1039–46
36. Guerrero-Bonilla L, Saldana D, Kumar V. Dense r-robust formations on lattices. In: *2020 IEEE international conference on robotics and automation (ICRA)*; 2020. pp. 6633–39
37. Stojanovic M. Recent advances in high-speed underwater acoustic communications. In: *IEEE Journal of Oceanic engineering*; 1996. vol. 23.2. pp. 125–36
38. Wang J. On sign-board based inter-robot communication in distributed robotic systems. In: *Robotics and automation*. 1045–50. *IEEE Comput. Soc. Press*. <https://doi.org/10.1109/ROBOT.1994.351219>
39. Cao YU, Fukunaga AS, Kahng A. Cooperative Mobile Robotics: Antecedents and Directions. In: *Autonomous Robots*; 1997. vol 4.1. <https://doi.org/10.1023/A:1008855018923>
40. IEEE I. IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. In: 1997. [Accessed 02 Jan 2022]. <https://doi.org/10.1109/IEEESTD.1997.85951>
41. Luu BB, et al. A soldier-robot ad hoc network. In: *Fifth Annual IEEE international conference on pervasive computing and communications workshops percom workshops 2007 19-23 March, 2007, White Plains, New York, USA*. IEEE Computer Society; 2007. pp. 558–563. <https://doi.org/10.1109/PERCOM.W.2007.13>
42. Houda K, Lakel R. Synchronized communication in a set of autonomous mobile robots using bluetooth technology. In: *Procedia Computer Science*; 2015. vol 73. pp. 154–61. 10.1016/j.procs.2015.12.061
43. Naik G, Choudhury B, Park J-M. IEEE 802.11bd & 5G NR V2X: Evolution of radio access technologies for V2X communications. In: *IEEE Access*; 2019. vol 7. pp. 70169–84. <https://doi.org/10.1109/ACCESS.2019.2919489>
44. Ertürk MA, et al. A Survey on LoRaWAN architecture, protocol and technologies. In: *Future internet*; 2019. vol. 11.10. pp. 216. <https://doi.org/10.3390/fi11100216>

45. Feng D, et al. Ultra-reliable and low-latency communications: applications, opportunities and challenges. In: Science China Information Sciences; 2021. vol. 64.2. pp. 1–12. <https://doi.org/10.1007/s11432-020-2852-1>
46. Mourad A, et al. A Baseline Roadmap for Advanced Wireless Research Beyond 5G. In: Electronics; 2020. vol. 9.2. <https://doi.org/10.3390/electronics9020351>
47. Van den Berg J, Lin M, Manocha D. Reciprocal velocity obstacles for real-time multi-agent navigation. In: IEEE International conference on robotics and automation. IEEE; 2008. pp. 1928–35
48. Koutsiamanis R-A, et al. From best effort to deterministic packet delivery for wireless industrial IoT networks. In: IEEE Transactions on industrial informatics; 2018. vol. 14.10. pp. 4468–80. <https://doi.org/10.1109/TII.2018.2856884>
49. Fan X, et al. UAV-assisted data dissemination in delay-constrained VANETs. In: Mobile information systems; 2018. vol. 2018. pp. 1–12. <https://doi.org/10.1155/2018/8548301>
50. Sarr C, et al. Bandwidth estimation for IEEE 802.11-based ad hoc networks. In: IEEE transactions on mobile computing; 2008. vol. 7.10. pp. 1228–41. <https://doi.org/10.1109/TMC.2008.41>
51. Yan C, et al. A Comprehensive survey on UAV communication channel modeling. In: IEEE Access; 2019. vol. 7. pp. 107769–92. <https://doi.org/10.1109/ACCESS.2019.2933173>
52. Khodayi-mehr R, Kantaros Y, Zavlanos MM. Distributed state estimation using intermittently connected robot networks. In: IEEE transactions on robotics; 2019. vol. 35.3. pp. 709–24. <https://doi.org/10.1109/TRO.2019.2897865>
53. Gupta L, Jain R, Vaszkun G. Survey of important issues in UAV communication networks. In: IEEE communications surveys & tutorials; 2016. vol. 18.2. pp. 1123–52. <https://doi.org/10.1109/COMST.2015.2495297>
54. Deng L, et al. Delay-constrained topology-transparent distributed scheduling for MANETs. In: IEEE transactions on vehicular technology; 2021. vol. 70.1. pp. 1083–88. <https://doi.org/10.1109/TVT.2020.3046856>
55. Elmezughi MK, Thomas A. Performance analysis of OFDM scheme with channel estimation and doppler shift effects. In: Albahit journal of applied sciences; 2021. vol. 2.1. pp. 2–8
56. IEEE I. P1920.2: Working group on vehicle-to-vehicle communications for unmanned aircraft systems; 2022
57. Ramphull D, et al. A review of mobile ad hoc NETWORK (MANET) protocols and their applications. In: 2021 5th International Conference on Intelligent Computing and Control Systems (ICICCS); 2021. pp. 204–11. <https://doi.org/10.1109/ICICCS51141.2021.9432258>
58. Al-Heety OS, et al. A comprehensive survey: benefits, services, recent works, challenges, security, and use cases for SDN-VANET. In: IEEE Access; 2020. vol. 8. pp. 91028–47. <https://doi.org/10.1109/ACCESS.2020.2992580>
59. Srivastava A, Prakash J. Future FANET with application and enabling techniques: Anatomization and sustainability issues. In: Computer science review; 2021. vol. 39. pp. 100359. <https://doi.org/10.1016/j.cosrev.2020.100359>
60. Tahir MN, Katz M. Performance evaluation of IEEE 802.11p, LTE and 5G in connected vehicles for cooperative awareness. In: Engineering Reports; 2021. <https://doi.org/10.1002/eng2.12467>
61. Hameed Mir Z, Filali F. LTE and IEEE 802.11p for vehicular networking: a performance evaluation. In: EURASIP journal on wireless communications and networking; 2014. vol. 2014.1. <https://doi.org/10.1186/1687-1499-2014-89>
62. Ziouva E, Antonakopoulos T. CSMA/CA performance under high traffic conditions: throughput and delay analysis. In: Computer communications; 2002. vol. 25.3. pp. 313–21. [https://doi.org/10.1016/S0140-3664\(01\)00369-3](https://doi.org/10.1016/S0140-3664(01)00369-3)
63. Adelantado F, et al. Understanding the limits of LoRaWAN. In: IEEE communications magazine; 2017. vol. 55.9. pp. 34–40. <https://doi.org/10.1109/MCOM.2017.1600613>
64. Shahin N, Ali R, Kim Y-T. Hybrid slotted-CSMA/CA-TDMA for efficient massive registration of IoT devices. In: IEEE Access; 2018. vol. 6. pp. 18366–82. <https://doi.org/10.1109/ACCESS.2018.2815990>
65. Petrosky EE, Michaels AJ, Ridge DB. Network Scalability Comparison of IEEE 802.15.4 and Receiver-Assigned CDMA. IEEE internet of things journal; 2019. vol. 4. pp. 6060–69. <https://doi.org/10.1109/JIOT.2018.2884455>
66. Shumeye Lakew D, et al. Routing in flying ad hoc networks: a comprehensive survey. In: IEEE communications surveys & tutorials; 2020. vol. 22.2. pp. 1071–120. 10.1109/COMST.2020.2982452. **This paper provides an overview of ad-hoc networks, their taxonomy, features (such as mobility models) and some of the protocols commonly employed. This holds relevance for ad-hoc networks in general.**
67. Hentati AI, Fourati LC. Comprehensive survey of UAVs communication networks. In: Computer Standards & Interfaces; 2020. vol. 72. pp. 103451. <https://doi.org/10.1016/j.csi.2020.103451>
68. Chataut R, Akl R. Massive MIMO systems for 5G and beyond networks-overview, recent trends, challenges, and future research direction. In: Sensors (Basel, Switzerland); 2020. vol. 20.10. <https://doi.org/10.3390/s20102753>
69. Liu X, Wang X. Efficient antenna selection and user scheduling in 5G Massive MIMO-NOMA system. In: IEEE 83rd Vehicular Technology Conference (VTC Spring). IEEE. 2016;2016:1–5. <https://doi.org/10.1109/VTCspring.2016.7504208>
70. Chen H, et al. Ultra-reliable low latency cellular networks: use cases, challenges and approaches. In: IEEE Communications Magazine; 2018. vol. 56.12. pp. 119–25. <https://doi.org/10.1109/MCOM.2018.1701178>
71. Zeng Y, Wu Q, Zhang R. Accessing from the sky: A tutorial on UAV communications for 5G and beyond. In: Proceedings of the IEEE; 2019. vol. 107.12. pp. 2327–75. <https://doi.org/10.1109/JPROC.2019.2952892>
72. Molina-Masegosa R, Gozalvez J. System level evaluation of LTE-V2V Mode 4 communications and its distributed scheduling. In: IEEE 85th vehicular technology conference (VTC Spring). IEEE. 2017;2017:1–5. <https://doi.org/10.1109/VTCspring.2017.8108463>
73. Miao L, Virtusio JJ, Hua K-L. PC5-based cellular-V2X evolution and deployment. In: Sensors (Basel, Switzerland); 2021. vol. 21.3. <https://doi.org/10.3390/s21030843>
74. Le T-K, Salim U, Kaltenberger F. An Overview of Physical Layer Design for Ultra-Reliable Low-Latency Communications in 3GPP Releases 15, 16, and 17. In: IEEE Access; 2021. vol. 9. pp. 433–44. <https://doi.org/10.1109/ACCESS.2020.3046773>
75. Ali R, et al. URLLC for 5G and Beyond: Requirements, enabling incumbent technologies and network intelligence. In: IEEE Access; 2021. vol. 9. pp. 67064–95. <https://doi.org/10.1109/ACCESS.2021.3073806>
76. Navarro-Ortiz J, et al. A Survey on 5G Usage scenarios and traffic models. In: IEEE Communications Surveys & Tutorials; 2020. vol. 22.2. pp. 905–29. <https://doi.org/10.1109/COMST.2020.2971781>
77. Lagen S, Patriciello N, Giupponi L. Cellular and Wi-Fi in unlicensed spectrum: competition leading to convergence. In: 2020 2nd 6G Wireless Summit (6G SUMMIT).[S.1]. IEEE; 2020. pp. 1–5. <https://doi.org/10.1109/6GSUMMIT49458.2020.9083786>
78. Di Zhang D, et al. Performance Analysis of FD-NOMA-Based Decentralized V2X Systems. In: IEEE transactions on communications; 2019. vol. 67.7. pp. 5024–36. <https://doi.org/10.1109/TCOMM.2019.2904499>

79. Shahab MB, et al. Grant-Free Non-Orthogonal Multiple Access for IoT: A Survey. *IEEE communications surveys & tutorials*; 2020. vol. 22.3. pp. 1805–38. <https://doi.org/10.1109/COMST.2020.2996032>
80. Soria E, Schiano F, Floreano D. Distributed predictive drone swarms in cluttered environments. In: *IEEE Robotics and Automation Letters*; 2021. vol. 7.1
81. Zhou X, et al. Decentralized spatial-temporal trajectory planning for multicopter swarms. In: [arXiv:2106.12481](https://arxiv.org/abs/2106.12481); 2021
82. Tordesillas J, How JP. MADER: Trajectory planner in multiagent and dynamic environments. In: *IEEE Transactions on Robotics*; 2021
83. Luis CE, Vukosavljev M, Schoellig AP. Online trajectory generation with distributed model predictive control for multi-robot motion planning. In: *IEEE Robotics and Automation Letters*; 2020. vol. 5.2. pp. 604–11
84. Kepler ME, Stilwell DJ. An approach to reduce communication for multi-agent mapping applications. In: 2020 IEEE/RSJ international conference on intelligent robots and systems (IROS); 2020. pp. 4814–20. <https://doi.org/10.1109/IROS45743.2020.9341117>
85. Trawny N, Roumeliotis SI, Giannakis GB. Cooperative multi-robot localization under communication constraints. In: 2009 IEEE international conference on robotics and automation; 2009. pp. 4394–400. <https://doi.org/10.1109/ROBOT.2009.5152606>
86. Muralidharan A, Mostofi Y. Communication-aware robotics: exploiting motion for communication. In: *Annual Review of Control, Robotics, and Autonomous Systems*; 2021. vol. 4.1. pp. 115–139. <https://doi.org/10.1146/annurev-contr-071420-080708>
87. Grancharova A, et al. UAVs trajectory planning by distributed MPC under radio communication path loss constraints. In: *Journal of Intelligent & Robotic Systems*; 2015. vol. 79.1. pp. 115–34
88. Zeng T, et al. Joint communication and control for wireless autonomous vehicular platoon systems. In: *IEEE Transactions on Communications*; 2019. <https://doi.org/10.1109/TCOMM.2019.2931583>
89. Zeng T, et al. Wireless communications and control for swarms of cellular-connected UAVs. In: 2018 52nd Asilomar Conference on Signals, Systems, and Computers; 2018. <https://doi.org/10.1109/ACSSC.2018.8645472>
90. Dimarogonas D, Johansson K. Bounded control of network connectivity in multi-agent systems. In: *IET Control Theory & Applications*; 2010. vol. 4.8. pp. 1330–38. <https://doi.org/10.1049/iet-cta.2009.0229>
91. Ji M, Egerstedt M. Distributed coordination control of multiagent systems while preserving connectedness. In: *IEEE Transactions on Robotics*; 2007. vol. 23.4. pp. 693–703. <https://doi.org/10.1109/TRO.2007.900638>
92. Stachura M, Frew EW. Cooperative target localization with a communication-aware unmanned aircraft system. In: *Journal of Guidance, Control, and Dynamics*; 2011. vol. 34.5. pp. 1352–62. <https://doi.org/10.2514/1.51591>
93. Olivieri de Souza BJ, Endler M. Coordinating movement within swarms of UAVs through mobile networks. In: 2015 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops); 2015. pp. 154–59. <https://doi.org/10.1109/PERCOMW.2015.7134011>
94. Mardani A, Chiaberge M, Giacccone P. Communication-aware UAV path planning. In: *IEEE Access*; 2019. vol. 7. pp. 52609–21. <https://doi.org/10.1109/ACCESS.2019.2911018>
95. Hsieh MA, et al. Maintaining network connectivity and performance in robot teams. In: *Journal of field robotics*; 2008. vol. 25. pp. 111–131
96. Schouwenaars T, et al. Multivehicle path planning for nonlinear-of-sight communication. In: *Journal of Field Robotics*; 2006. vol. 23. pp. 269–90. <https://doi.org/10.1002/rob.20119>
97. Fridman A, et al. Distributed path planning for connectivity under uncertainty by ant colony optimization. In: 2008 American Control Conference; 2008. pp. 1952–58. <https://doi.org/10.1109/ACC.2008.4586778>
98. Hayat S, et al. Multi-objective UAV path planning for search and rescue. In: 2017 IEEE International Conference on Robotics and Automation (ICRA); 2017. pp. 5569–74. <https://doi.org/10.1109/ICRA.2017.7989656>
99. Doniec M, et al. Using optical communication for remote underwater robot operation. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE; 2010. pp. 4017–4022
100. Vasilescu I, et al. Data collection, storage, and retrieval with an underwater sensor network. In: Proceedings of the 3rd international conference on embedded networked sensor systems. San Diego, California, USA. Association for Computing Machinery; 2005. pp. 154–65. <https://doi.org/10.1145/1098918.1098936>
101. Zolich A, et al. Survey on communication and networks for autonomous marine systems. In: *Journal of Intelligent & Robotic Systems*; 2019. vol. 95.3. pp. 789–813
102. Hollinger GA, et al. Communication protocols for underwater data collection using a robotic sensor network. In: 2011 IEEE GLOBECOM Workshops (GC Wkshps); 2011. pp. 1308–13. <https://doi.org/10.1109/GLOCOMW.2011.6162397>
103. Anderson J, Hollinger GA. Communication planning for cooperative terrain-based underwater localization. In: *Sensors*; 2021. vol. 21.5. p. 1675. <https://doi.org/10.3390/s21051675>
104. Best G, et al. Planning-aware communication for decentralized multi-robot coordination. In: 2018 IEEE International Conference on Robotics and Automation (ICRA); 2018. pp. 1050–57. <https://doi.org/10.1109/ICRA.2018.8460617>
105. Alshehri A, Miller T, Sonenberg L. Modeling communication of collaborative multiagent system under epistemic planning. In: *International Journal of Intelligent Systems*; 2021. vol. 36.10. pp. 5959–80. <https://doi.org/10.1002/int.22536>
106. Unhelkar VV, Shah JA. Contact: Deciding to communicate during time-critical collaborative tasks in unknown, deterministic domains. In: Thirtieth AAAI Conference on Artificial Intelligence; 2016
107. Tsiogkas N, Lane DM. Towards an online approach for knowledge communication planning: Extended Abstract. In: 2019 international symposium on multi-robot and multi-agent systems (MRS); 2019. pp. 142–44. <https://doi.org/10.1109/MRS.2019.8901086>
108. Marcotte RJ, et al. Optimizing multi-robot communication under bandwidth constraints. In: *Autonomous Robots*; 2020. vol. 44.1. pp. 43–55
109. Otsu K. Supervised autonomy for communication-degraded subterranean exploration by a robot team. In, et al. IEEE Aerospace Conference. IEEE. 2020;2020:1–9. <https://doi.org/10.1109/AERO47225.2020.9172537>.
110. Tranzatto M. CERBERUS in the DARPA Subterranean Challenge. In: *Science Robotics*; 2022. vol. 7.66. p. eabp9742
111. Mascarich F. A self-deployed multi-channel wireless communications system for subterranean robots. In, et al. IEEE Aerospace Conference. IEEE. 2020;2020:1–8. <https://doi.org/10.1109/AERO47225.2020.9172496>.
112. Kantaros Y, Zavlanos MM. Distributed communication-aware coverage control by mobile sensor networks. In: *Automatica*; 2016. vol. 63. pp. 209–20. <https://doi.org/10.1016/j.automatica.2015.10.035>

113. Zavlanos MM, Ribeiro A, Pappas GJ. Network integrity in mobile robotic networks. In: IEEE Transactions on Automatic Control; 2013. vol. 58.1. pp. 3–18. <https://doi.org/10.1109/TAC.2012.2203215>
114. Kassir A, Fitch R, Sukkarieh S. Decentralised information gathering with communication costs. In: 2012 IEEE international conference on robotics and automation; 2012. pp. 2427–32. <https://doi.org/10.1109/ICRA.2012.6224806>
115. Fink J, Ribeiro A, Kumar V. Robust control for mobility and wireless communication in Cyber-Physical systems with application to robot teams. In: Proceedings of the IEEE; 2012. vol. 100.1. pp. 164–78. <https://doi.org/10.1109/JPROC.2011.2161427>
116. Le Ny J, Ribeiro A, Pappas GJ. Adaptive communication-constrained deployment of mobile robotic networks. In: American control conference (ACC). Montreal, QC. IEEE; 2012;2012:3742–7. <https://doi.org/10.1109/ACC.2012.6314804>
117. Stephan J, et al. Concurrent control of mobility and communication in multirobot systems. In: IEEE Transactions on Robotics; 2017. vol. 33.5. pp. 1248–54. <https://doi.org/10.1109/TRO.2017.2705119>
118. Halsted T, et al. A survey of distributed optimization methods for multi-robot systems; 2021. [arXiv:2103.12840](https://arxiv.org/abs/2103.12840) [cs]
119. Amato C, et al. Decentralized control of partially observable Markov decision processes. In: 52nd IEEE conference on decision and control. IEEE; 2013. pp. 2398–405
120. Parker L. ALLIANCE: an architecture for fault tolerant multirobot cooperation. In: IEEE Transactions on Robotics and Automation; 1998. vol. 14.2. pp. 220–40. <https://doi.org/10.1109/70.681242>
121. Mammeri Z. Reinforcement learning based routing in networks: Review and classification of approaches. In: IEEE Access; 2019. vol. 2019. pp. 55916–50. <https://doi.org/10.1109/ACCESS.2019.2913776>
122. Zheng Z, Sangaiah AK, Wang T. Adaptive communication protocols in flying ad hoc network. In: IEEE Communications Magazine; 2018. vol. 56.1. pp. 136–42. <https://doi.org/10.1109/MCOM.2017.1700323>
123. Bithas PS, et al. A Survey on Machine-Learning Techniques for UAV-Based Communications. In: Sensors (Basel, Switzerland); 2019. vol. 19.23. <https://doi.org/10.3390/s19235170>
124. Wang J-L, et al. Machine learning based rapid 3D channel modeling for UAV communication networks. In: CCNC 2019: 2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC), Piscataway, NJ. IEEE; 2019. pp. 1–5. <https://doi.org/10.1109/CCNC.2019.8651718>
125. Rajeswaran A, et al. Towards generalization and simplicity in continuous control. In: Proceedings of the 31st international conference on neural information processing systems, Long Beach, California, USA. Curran Associates Inc; 2017. pp. 6553–64
126. Tobin J, et al. Domain randomization for transferring deep neural networks from simulation to the real world. In: IEEE/RSJ International Conference on Intelligent Robots and Systems; 2017. pp. 23–30. <https://doi.org/10.1109/IROS.2017.8202133>
127. Everett M, Chen YF, How JP. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In: IEEE/RSJ international conference on intelligent robots and systems; 2018. pp. 3052–59. <https://doi.org/10.1109/IROS.2018.8593871>
128. Sartoretti G, et al. PRIMAL: pathfinding via reinforcement and imitation multi-agent learning. In: IEEE Robotics and Automation Letters; 2019. vol. 4.3. pp. 2378–85
129. Prorok A, et al. The holy grail of multi-robot planning: Learning to generate online-scalable solutions from offline-optimal experts In: Autonomous Agents and Multi-Agent Systems (AAMAS); 2022
130. Paulos J, et al. Decentralization of multiagent policies by learning what to communicate. In: 2019 International Conference on Robotics and Automation (ICRA). IEEE; 2019. pp. 7990–96
131. Scarselli F, et al. The graph neural network model. In: IEEE Transactions on Neural Networks; 2009. vol. 20.1. pp. 61–80
132. Gama F, et al. Convolutional neural network architectures for signals supported on graphs. In: IEEE Trans. Signal Process; 2019. vol. 67.4. pp. 1034–49
133. Prorok A. Graph neural networks for learning robot team coordination. In: Federated AI for Robotics Workshop, IJCAI-ECAI/ICML/AAMAS 2018; [arXiv:1805.03737](https://arxiv.org/abs/1805.03737) [cs]; 2018
134. Kortvelesy R, Prorok A. ModGNN: Expert policy approximation in multi-agent systems with a modular graph neural network architecture. In: International Conference on Robotics and Automation (ICRA); 2021
135. Khan A, et al. Graph policy gradients for large scale robot control. Ed. Kaelbling LP, Kragic D, Sugiura K. Proceedings of the conference on robot learning. PMLR; 2020. vol. 100. pp. 823–34
136. Tolstaya E, et al. Learning decentralized controllers for robot swarms with graph neural networks; 2019. [arXiv:1903.10527](https://arxiv.org/abs/1903.10527) [cs]
137. Li Q, et al. Graph neural networks for decentralized multi-robot path planning. In: Autonomous Agents and Multi-Agent Systems (AAMAS); 2020
138. Hu T-K, et al. Scalable perception-action-communication loops with convolutional and graph neural networks. In: IEEE Transactions on Signal and Information Processing over Networks; 2022. vol. 8. pp. 12–24. <https://doi.org/10.1109/TSIPN.2021.3139336>
139. Li Q, et al. Message-aware graph attention networks for large-scale multi-robot path planning. In: IEEE Robotics and Automation Letters; 2021. vol. 6.3. pp. 5533–40
140. Sukhbaatar S, Szlam A, Fergus R. Learning multiagent communication with backpropagation. In: Proceedings of the 30th international conference on neural information processing systems, '16. Curran Associates Inc; 2016. pp. 2252–60
141. Singh A, Jain T, Sukhbaatar S. Learning when to communicate at scale in multiagent cooperative and competitive tasks. In: International Conference on Learning Representations; 2018
142. Jiang J, Lu Z. Learning attentional communication for multi-agent cooperation. In: Proceedings of the 32nd international conference on neural information processing systems, '18. Red Hook, NY, USA. Curran Associates Inc; 2018. pp. 7265–75
143. Das A, et al. TarMAC: Targeted multi-agent communication. In: Proceedings of the 36th International Conference on Machine Learning; 2019. pp. 1538–46. PMLR. **This paper demonstrates a learning-based approach to multi-robot communications, particularly, learning what and when to communicate so as to maximize task-specific rewards.**
144. Serra-Gómez, Á, et al. With whom to communicate: learning efficient communication for multi-robot collision avoidance. In: 2020 IEEE/RSJ international conference on intelligent robots and systems (IROS); 2020. <https://doi.org/10.1109/IROS45743.2020.9341762>
145. Blumenkamp J, Prorok A. The emergence of adversarial communication in multi-agent reinforcement learning; 2020. [arXiv:2008.02616](https://arxiv.org/abs/2008.02616) [cs]
146. Mechraoui A, et al. Co-design for wireless networked control of an intelligent mobile robot. In: ICINCO 2009-6th International Conference on Informatics in Control, Automation and Robotics; 2009, p. 7
147. Mechraoui A, et al. A co-design distributed Bayesian approach for decision and scheduling of WNCs. In: IFAC Proceedings Volumes; 2011. vol. 44.1. pp. 14970–75
148. Blumenkamp J, et al. A framework for real-world multi-robot systems running decentralized GNN-based policies. In: International Conference on Robotics and Automation (ICRA); 2022