Social-Aware Coordination of Multi-Robot Systems Based on Institutions

José N. Pereira, Porfírio Silva, Pedro U. Lima and Alcherio Martinoli

Abstract Institutional robotics (IR) is an approach to the coordination of multi-robot systems that draws inspiration from social sciences, namely from institutional economics. Using the concept of institution, it aims to provide a comprehensive strategy for specifying social interactions (e.g., norms, roles, hierarchies) among robots. In previous work, we have introduced a control methodology for multi-robot systems that takes into account institutions in order to create an Institutional Agent Controller (IAC) that captures such social interactions. In this chapter, the IAC design methodology is validated in a case study concerned with a swarm of 40 real, resource-constrained robots which has to maintain wireless connectivity. We then investigate a second case study dealing with more complex social interactions, showing that institutional roles can effectively help a multi-robot system to coordinate and improve performance in a given task of social nature. Given the fact that institutions are one of the tools in use within human societies to shape social interactions, our intuition is that IR can also facilitate coordination with humans in scenarios involving many-to-many human-robot interactions. We discuss how the IR concepts and the IAC design methodology can be implemented in real-world scenarios where multiple robots must interact with multiple humans in a socially-aware manner.

José N. Pereira and Alcherio Martinoli
Distributed Intelligent Systems and Algorithms Laboratory - EPFL, EPFL ENAC IIE DISAL, Station 2, CH-1015 Lausanne, Switzerland, e-mail: `{jose.pereira,alcherio.martinoli}@epfl.ch`

Porfírio Silva and Pedro U. Lima
Institute for Systems and Robotics - Instituto Superior Técnico, Universidade de Lisboa, Av. Rovisco Pais 1, 1049 Lisboa, Portugal, e-mail: `{porfiriosilva,pal}@isr.ist.utl.pt`

This work was partially supported by Fundação para a Ciência e a Tecnologia (FCT) through grants SFRH/BD/33671/2009 (first author, as part of the Joint Doctoral Program IST-EPFL) and SFRH/BPD/35862/2007 (second author), as well as by FCT ISR/IST Pluriannual funding through the PIDDAC program funds, and by EU under FP7/2007-2013 - Challenge 2 - Cognitive Systems, Interaction, Robotics - grant agreement 601033 - MOnarCH.

1 Introduction

In robotics, the transition from constrained laboratory environments to real-world environments is not a trivial step. For human actors carrying out any task in real-world environments, even without any robotic systems present, the need for coordination with other actors is ubiquitous. Consider that simply traveling from point *A* to point *B* requires coordinating with others, be it on foot (using common sense rules to avoid colliding with others), by car (using a more elaborate set of rules, i.e., the “road code”), or any other mode of transport. More importantly, most tasks in our daily lives involve the trade of services or goods for money. This requires not only coordinating with others in order to physically carry out such a trade but also the underlying acceptance that the money being received can then be accepted by some other actors. The collective acceptance of such an idea is in itself a form of society-wide coordination.

For multi-robot systems to be truly immersed in real-world environments this type of coordination must be considered. Robots need to consider complex social interactions with multiple anonymous robots and multiple anonymous human actors. This anonymity reflects not only the fact that interacting agents might never have met before but also that such agents might not have to coexist at the same time or place for the interaction to occur.

Our goal is to formalize rules specifying social interactions for specific tasks in a way that resembles the organizational aspects of human society. Our intuition is that by doing so we will be able to consider complex social interactions within multi-robot systems, ease the effort of their transition to real-world environments populated with human actors, and facilitate coordination with such actors in scenarios involving many-to-many human-robot interactions.

To do so, we follow the *Institutional Robotics* (IR) approach [23] to the coordination of multi-robot systems (described in Section 2). This approach takes inspiration from social sciences, namely from *institutional economics* [10], and aims to provide a comprehensive strategy for specifying complex social interactions among a team of robots and possibly between a team of robots and human actors. In [20], we formalized institutions - the central concept in IR - using an abstract representation (executable Petri Nets), allowing their design and execution for multi-robot systems, so as to obtain behaviors capturing the social interactions of interest. Our methodology (described in Section 3) composes a set of institutions in order to create an institutional agent controller able to execute a desired task and observe the specified social interactions.

An initial validation study was performed in [20], by comparing our methodology with other approaches. To do so, we considered a swarm robotics case study concerned with a robot swarm which has to maintain wireless connectivity and a certain degree of spatial compactness, taking into account only simple social interactions. However, the validation presented in [20] was carried out only in simulation. In Section 4, we advance our validation effort by considering a real-world implementation of the case study with a swarm of up to 40 real, resource-constrained robots, further increasing our confidence in the approach's robustness and scalability. This real-world implementation was reported in [19].

Considering a case study concerned only with simple interactions allows us to perform a more grounded validation effort. However, we are interested in tackling more complex social interactions with the IR approach. In Section 5, we present a second case study (briefly described in [18]) where a team of robots must coordinate while navigating through the environment in order to accomplish a transportation task. In this case study we increase the complexity of the social interactions by focusing on a specific form of institution: the institutional role. We compare the IR approach with a self-organized approach in order to identify situations in which it might prove advantageous.

The IR approach will also be considered in social robotics scenarios implemented in real-world human-populated environments. In Section 6, we discuss how the IR approach can enforce social-aware coordination in such a scenario, and how we will be able to verify our intuition about the impact of IR in many-to-many human-robot interactions.

2 Institutional Robotics and Related Work

Institutional robotics [23] is an approach to the coordination of multi-robot systems that draws inspiration from the social sciences, namely from institutional economics' concepts [10]. It combines the notions of institution [21, 11], coordination artifact [25], and environment [28], aiming to provide a comprehensive strategy for specifying social interactions (e.g., norms, roles, hierarchies) among robots. Under IR, robots are situated not only in a physical but also in an institutional environment, where their interactions are guided by institutions. Cooperation is achieved by this regulation of social interactions since the robots know not only how to behave in a given scenario but also what to expect from other robots and the environment.

In IR, the coordination system is a network of institutions. Institutions are coordination artifacts of different types (organizations, norms, hierarchies, social roles, etc). They are generic, meaning that they are not designed for any specific set of robots. Robots are able to modify, at some extent, not only the physical environment but also the institutional environment. From an institutional perspective, institutions are taken as the main tool of any sophisticated society, and individuals are both constructive within and constructed through institutional environments.

Market-based multi-robot coordination [7, 31] is a previous example of importing some economic views into robotics. Inspired by market mechanisms, researchers have proposed systems like MURDOCH [9] and TraderBots [8] to achieve flexible allocation of subtasks using auctions between robots. In these systems, robots act as agents trying to maximize their individual profits, calculated based on rewards from tasks and resources expended. The underlying assumption is that with every robot trying to maximize its individual profit, team coordination and efficiency will be improved. A limitation of the market-based approach is that, despite some application to the allocation of roles [26], the great majority of the work available only deals with task allocation, leaving other mechanisms (e.g., cooperative decision-making) out of the picture.

Self-organization is another possible, scalable mechanism that has been proposed for the coordination of, often large, distributed robotic systems [2, 1, 6]. Self-organizing systems are characterized as being fully reactive and relying on local interactions (and possibly local, broadcast communication), both between robots and between robots and the environment, in order to achieve coordination. However, the design of truly social distributed robotic systems should take into account the objective social interactions arising from the combination (and dependencies) of goals of heterogenous agents [5]. The simple, local interactions considered by self-organizing systems might not capture well this aspect.

3 Institutional Agent Controllers

We model institutions using a formal representation, leading to a standard design and execution platform (in real robots, submicroscopic realistic simulations, and microscopic multi-agent systems). Institutions encapsulate relevant behavioral rules for robots, specifying social interactions of different types among actors in a given scenario. They represent the basic building blocks for creating shared coordinated working environments. Moreover, concurrent execution of institutions has to be regulated since not all behaviors can be executed simultaneously. We use *Petri Nets* (PNs) as the formal framework and follow their usual definition as described in [3].

Formalizing institutions for modeling and execution of robot controllers means that we need to take into account robot's actions and sensor readings. *Executable Petri Nets* (EPNs) are PNs that have actions and boolean conditions (verifiable by sensor readings) associated with places and transitions, respectively. The basic intuition behind this definition is that by associating actions with places we are able to define which actions are to be executed at each time step. This is done simply by checking if the corresponding place is marked. By associating transitions with conditions verified by sensor readings we trigger state changes in the EPN due to changes in the robot's environment.

We represent each institution by an EPN that can be executed independently or together with other institutions. We also represent robot's *individual behaviors* by EPNs. While the institutions specify behaviors that have a *social nature*, i.e., they relate the robot to other robots in some way, the individual behaviors specify a set of basic behaviors that have exclusively an *individual nature*, i.e., they relate the robot with the surrounding environment and its own goals. The composition of the individual behavior with a set of institutions generates a robot controller.

Definition: An *Institution I* is a four-tuple $(Inst, initial_I, final_I, d_I)$ where:

- $Inst$ is an EPN;
- $initial_I, final_I \in Cdt$ are initial and final conditions for the execution of $Inst$;

- $d_I \in D = \{AllowAll, StopInd, StopInst, StopAll\}$ is the associated deontic operator.

The EPN *Inst* specifies the desired behavior that should be performed by the robot. This behavior is not always being executed, its start and end are dictated by conditions $initial_I$ and $final_I$, which the robot verifies at each time step. Thus, we say that an institution I at each time step can be *active* or *idle*. Each institution also includes a deontic operator d_I which is used when combining it with the robot's individual behavior and further institutions, allowing or stopping the concurrent execution of institutions and/or individual behavior. *Inst* must be designed, but institutions can be kept simple and further behavioral complexity is the result of composition, in a modular fashion.

EPNs can be represented by macro places in a hierarchical fashion, using two distinct layers. We consider that each institution I is part of a lower layer and is represented by one macro place m_I in the higher layer. By adding bidirectional arcs between each transition in I and m_I, we guarantee that if m_I is marked, I is active, otherwise it is idle. This allows us to compose our institutions at the higher layer where relationships among the institutions and the individual behavior should be specified while keeping relationships between actions and conditions separated in the lower layer.

The composition of individual behaviors and institutions is performed algorithmically by adding, in the higher layer, places and transitions that restrict their concurrent execution, according to the specification provided by the deontic operators. Both layers can be then merged algorithmically to obtain a full EPN that can be used as controller. This EPN is designated as the *Institutional Agent Controller* (IAC). Each robot runs its IAC in a social collective setting mediated by institutions. An example of a specific IAC for the wireless connected swarm case study is displayed in Fig. 2 and will be discussed in detail in the next section.

4 Validation of IAC methodology

In order to validate the IAC methodology for design and execution of robotic controllers using institutions we follow a two-phase approach. The first phase is to compare the IAC methodology with other methods, and assess its capability to replicate results obtained with such methods. In [20] we considered the wireless connected swarm case study, previously investigated in [16, 30, 29], and applied to it our IAC methodology. We performed submicroscopic simulations, characterized by a low degree of abstraction and where intra-robot details such as individual sensors and actuators are captured, both with the IAC designed and with a Finite State Automata (FSA) approach originally proposed in [16]. We concluded that the IAC methodology was able to replicate results obtained with the more traditional FSA approach.

The second phase is to validate the IAC methodology in real world environments. In this section we focus on this phase. To do so, we perform real world experiments of the case study and compare results obtained in reality with those obtained in submicroscopic simulations.

In [30] the authors report an implementation using a small number of real robots (4-8 robots) where local communication was achieved with a combination of a global wireless network and an overhead camera delimiting the communication range. A first goal of this work is to move one step further the realism of the physical implementation by using real local communication channels and a large number of robots (tests were performed with sets of 20 and 40 robots). We chose to use 40 robots in order to maintain as much as possible a parallel with the original case study experiments, where 40 simulated agents were used [29]. A second goal of such implementation is to show that the IR approach is able to handle such real scenarios, and in particular maintain the wireless connectivity of a swarm of 40 real, resource-constrained robots, further increasing our confidence in the approach's robustness and scalability.

4.1 Materials and Methods

Our platform is the *e-puck* robot [15], a differential drive robot of 7 cm in diameter. In order to endow the robots with scalable wireless communication capabilities, we use a radio communication module developed at DISAL [4]. This module is ZigBee-compliant and uses TinyOS [12]. A bounded communication range is obtained using software-controllable power emission and a dedicated hardware attenuator.

For implementing our submicroscopic simulations we used *Webots* [14], a flexible, 3D realistic simulator, and considered kinematic models of the *e-puck* robot. The original case study considered a perfect circular bounded communication radius and perfect package reception inside that radius (radial disk model). In this work, communication between *e-pucks* is also simulated realistically using the network simulation engine OMNeT++ [27] as a plugin for *Webots*. The OMNet++ engine handles channel coding, noise, fading signal propagation, as well as a non-circular communication footprint. Fig. 1-(a) offers a visualization of the *Webots* submicroscopic simulations. Fig. 1-(b) displays an image of the arena during execution taken with the overhead camera.

4.2 Task Description & Decentralized Control Algorithm

In the wireless connected swarm case study a decentralized control algorithm is implemented to maintain wireless connectivity and a certain degree of spatial compactness of a robotic swarm (with N robots) in an unbounded arena using exclusively, as information at the robot level, the current number of wireless connections to the neighbors. The communication is local and its bounded range is a parameter of the robotic system. Let X be the number of connections perceived by a robot.

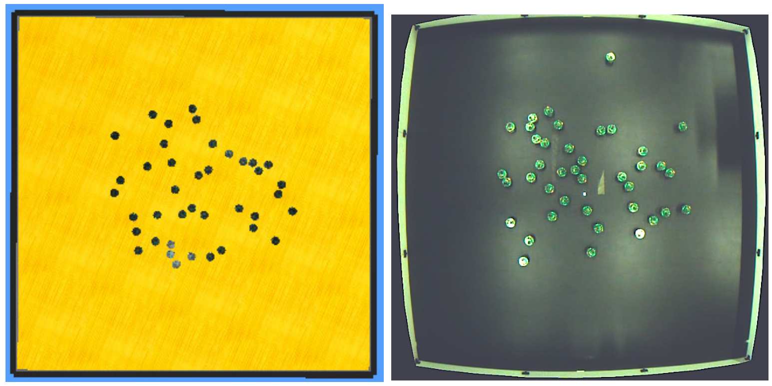

(a) (b)

Fig. 1 (a) *Webots* simulation screenshot, 40 *e-puck* robots simulated. (b) Real world experiment screenshot, 40 *e-puck* robots.

In the default state (defined as $forward$), the robot simply moves forward. If at any time the robot senses the loss of a connection and X falls below a threshold α (where $\alpha \in \{0,\ldots,N-1\}$), the robot assumes it is going in the wrong direction and switches to state *coherence*. In this state the robot performs a 180° turn in order to recover the lost connection. Upon recovering the lost connection, the robot performs a random turn and moves back to the default state. If the connection is not recovered, the robot simply moves to the default state. If an obstacle is detected the robot immediately switches to state *avoid*, where it performs obstacle avoidance for a given number of time steps, after which it returns to its previous state.

While this simple algorithm has limited robustness, it allows the swarm to maintain its connectivity to a certain extent, with its spatial compactness being controlled by the communication range and by the threshold α. It is implemented in [29] using a FSA controller with states defined as above.

4.3 Institutional Agent Controller

In our IAC implementation, robots execute an individual behavior *IndAv* (*Individual Avoidance*) and two institutions $T180$ (*Turn 180 degrees*) and TR (*Turn Random*), all specified by EPNs shown in the lower layer of Fig. 2. Individual behavior *IndAv* specifies a behavior relating the robot to its environment, consisting on simple obstacle avoidance. Institutions $T180$ and TR implement the social rules, dealing with loss and recovery of connections. $T180$ specifies that upon losing a connection the robot performs a 180° turn followed by moving forward for a small number of steps. Institution TR specifies that if a connection is recovered the robot performs a random degree turn.

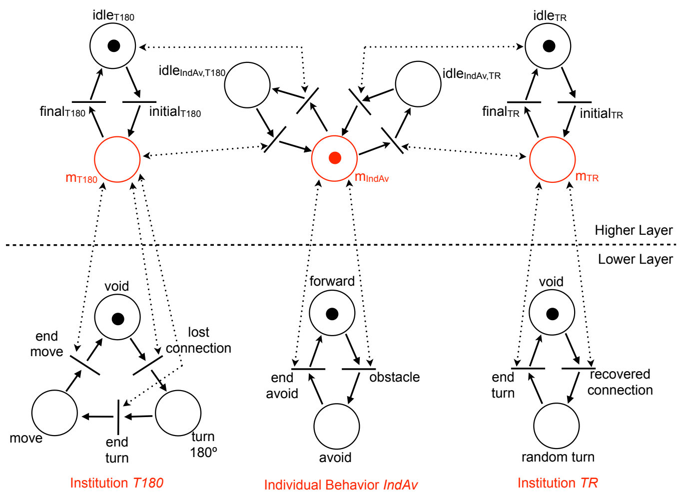

Fig. 2 IAC for the wireless connected swarm. Dotted arcs represent bidirectional arcs. Lower layer: EPNs for individual behavior *IndAv* and institutions $T180$ and TR. Higher layer: composition of individual behavior and institutions.

To consider institutions as defined in Section 3, we need initial and final conditions and deontic operators. For institution $T180$ we say that initial condition $initial_{T180}$ is “loss of connection detected and number of connections is less than α” and the final condition $final_{T180}$ is “move forward procedure has ended”. For institution TR we say that initial condition $initial_{TR}$ is “recovery of connection detected and previous number of connections is less than α” and the final condition $final_{TR}$ is “random turn procedure has ended”. The deontic operator associated with both institutions is $StopInd$, specifying that institutions and individual behavior cannot be executed concurrently.

We now have all the elements needed to obtain the IAC that specifies our desired behavior. The composition of the individual behavior *IndAv* and institutions $T180$ and TR (specified separately by EPNs shown in the lower layer of Fig. 2) is shown in the higher layer of Fig. 2. The final controller is the full EPN of Fig. 2, obtained after merging the two layers.

4.4 Experimental Setup

We replicated, to the best possible extent, the conditions of the original case study presented in [29]. Therein, the authors considered 40 robots in an unbounded arena performing the task over 10 000 seconds. In this work, we carry out experiments (both real robot experiments and *Webots* simulations) with sets of $N = 20$ and $N = 40$ robots in a 3 by 3 meters bounded arena performing the task over 1800 seconds. The connection threshold is dependent on the size of N and is set to $\alpha = 8$ for $N = 20$ and $\alpha = 16$ for $N = 40$. The communication radius of the *e-puck* is intended to be 0.7 meters, instead of the original 2.0 meters, in order to keep the ratio between communication and physical radius presented in the original paper. We set the transmission power of the *e-puck* communication module to an appropriate value that allows us to roughly achieve the desired communication radius.

To compare the performance of our submicroscopic simulations and real world experiments we performed 100 runs of the simulation for each $N = 20$ and $N = 40$, and 10 runs of real world experiments for $N = 20$ and 5 runs for $N = 40$. During runs we stored the number of time steps robots spent with each number of connections (between 0 and $N - 1$). We also recorded videos of the arena during the real world experiments using an overhead camera and the *SwisTrack* software [13]. We processed the videos offline, using *SwisTrack* to perform background subtractions and blob detection, in order to extract and store the position of each robot in each frame. We also stored information about the position of robots at each time step of our simulations.

4.5 Results and Discussion

In this work, we are interested in two main metrics that represent and allow us to analyze different aspects of the swarm behavior: connectivity and displacement.

Connectivity tells us, on average, how many robots have a particular number of wireless connections during the time needed to perform a run of the experiment. To measure connectivity we use data gathered by the robots about the number of time steps spent with each number of connections. Robots with α or more connections are not concerned with recovering lost connections and are likely to be moving away from the swarm. On the other hand, robots with less than α connections are actively trying to regain connections and are likely to be moving towards the swarm. Thus, we can expect the swarm connectivity to peak at α, i.e., at each time step we will have more robots with α connections than with any other number of connections.

Displacement measures the distance between the swarm center of mass and the center of the arena. Given the stochastic nature of the movement of the robots, displacement will start close to zero (runs start with robots gathered closely in the center of the arena) and will increase throughout the run. The motion of the swarm as a whole resembles a random walk through the arena. Comparing the values for displacement in tests made with simulated and real robots will give us additional insights into how well the simulator is capturing reality. This metric would be somewhat different if considered in the original case study, given that an unbounded arena was considered.

In Fig. 3-(a) and Fig. 3-(b) we present the connectivity metric results for $N = 20$ and $N = 40$. In green we display results obtained with submicroscopic simulations, while in red and blue we display results with real robots. The blue line was obtained with data about number of connections as perceived and recorded by the robots. On

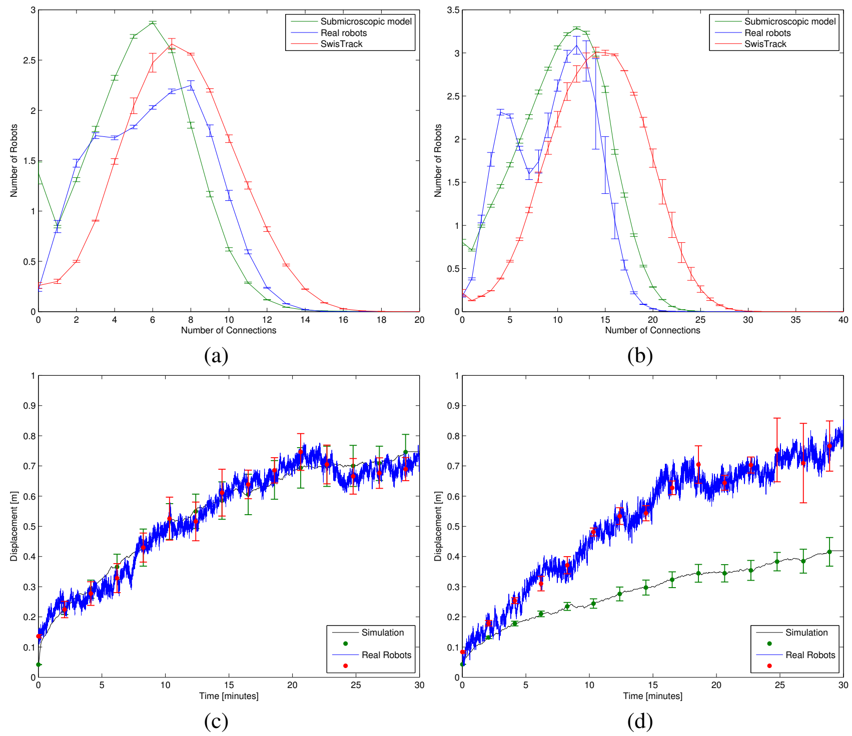

Fig. 3 (Top) Connectivity metric: average number of robots with a particular number of connections during a run. (Bottom) Displacement metric: average distance of swarm center of mass to arena center throughout a run. (Left) Results for 20 robots and $\alpha = 8$. (Right) Results for 40 robots and $\alpha = 16$. Variance shown for different runs.

the other hand, the red line was obtained in offline processing using *SwisTrack* by counting, for each robot, how many other robots were present in a 0.7 meters range, somehow emulating a perfectly radial communication disk. The differences in these two lines can be explained by the spatially irregular coverage of the wireless radio communications. The blue line reflects more accurately this noisy nature by spreading the number of robots more evenly between 3 and 9 connections in Fig. 3-(a) and producing a second local maximum for 4 connections in Fig. 3-(b). This maximum can be explained by the increase in N and α. The increase in α forces robots to try to keep more neighbors in their communication radius, leading to robots aggregating in a smaller space. This effect is magnified by the increase of robots in the swarm. Thus, when robots lose or gain connections they lose 1 or 2 connections with $N = 20$ but they lose 4 or 5 connections with $N = 40$. The video data processed with *SwisTrack* always gives the correct number of neighbors since all robot positions are known, thus the red line better reflects the overall swarm behavior. We can see that connectivity measured with *SwisTrack* has a very good agreement with the connectivity measurements obtained in our submicroscopic simulations. The slight shift of the curve of the simulations in relation to the curve of *SwisTrack*, representing that robots have on average slightly less connections, is most likely a product of the inclusion of wireless communication realism (noisy fading and ellipsoidal communication area) in the simulations through the OmNET++ plugin. These results show that, despite the high influence of noise in real world wireless communication, the overall swarm behavior implemented using an IR distributed control approach is able to maintain the expected connectivity. The results also show a very good agreement with the results presented in the original case study work [29].

In Fig. 3-(c) and Fig. 3-(d) we present the displacement metric results for $N = 20$ and $N = 40$. Real robots results are obtained only using the video data processed with *SwisTrack*, since robots do not have localization capabilities and are unaware of their own location as well as the location of others. As expected, displacement distance is close to zero at the beginning and increases throughout the run. For $N = 20$, submicroscopic simulations and real robot experiments show perfect agreement. However, for $N = 40$, despite distance increasing in both simulation and real robots, we observe that the rate of increase is doubled from simulation to real robots. A possible explanation for this effect is the difference in the obstacle avoidance behavior. While in submicroscopic simulations *e-pucks* are considered as perfect cylindrical blocks, in reality *e-pucks*' bodies are translucent. This leads to some collisions between robots, being this effect greatly increased when the number of robots is doubled and they are forced to aggregate in a smaller space (because α is also doubled). Robots motion becomes less predictable and more stochastic and as a result the displacement of the whole swarm is increased, much in the same manner as a random walk with increased turning probability.

5 Coordination Through Institutional Roles

As discussed previously, institutions can take several forms: organizations, norms, hierarchies, roles, etc. In the wireless connected swarm case study, we experimented with institutional norms in a low complexity task, used for IAC validation only. We now focus on another important institutional form, the institutional role. Our objective is to observe if coordination of a multi-robot system can be improved by using institutional roles, possibly in combination with other types of institutions. Also, the task to be accomplished by the robots in this case study is of higher complexity (w.r.t. the wireless connected swarm case study), with complexity of social interactions among robots depending on the approach taken.

We consider the following scenario and task. Robots are situated in an arena consisting of two rooms connected by a narrow corridor (see Fig. 4-(a)). The width of the corridor allows only for one robot. In the left room robots can pick up "virtual payloads" (in infinite supply) that can be deployed in the right room. Both picking up and deploying virtual payloads happen after a fixed amount of time has elapsed since the robots enter the respective rooms. In order for the robots to recognize their topological location, the walls have different colors, yellow in the left room, green in the right room and blue in the corridor.

The team goal is to maximize the number of deployed virtual payloads. Robots pick up the virtual payload in the left room. They must then navigate through the corridor and deploy the payload in the right room. The corridor connecting the rooms is too narrow for two robots moving in opposite directions to pass one another. In order to avoid congestion in the corridor, the traffic between the two rooms must to be coordinated so that robots only attempt to traverse the corridor in one direction at a time.

We show how such coordination can be achieved when the constituent robots are given the capacity to assume an institutional role, that of traffic regulator. We compare this institutional approach to a self-organized approach to the same task and try to identify in which situations one is preferable to the other.

5.1 Institutional Approach

We designate by *transporting* robots all robots that are transporting virtual payloads, and thus, actively accomplishing the task. Robots performing an institutional role are designated as *regulators* (we use also *traffic regulators* interchangeably).

5.1.1 Transporting Robots

Initially, all robots are transporting robots. They are placed randomly in the two rooms (see Fig. 4-(a)) where they attempt to locate a wall and perform a wall-following behavior by keeping a wall on their right hand side, using readings from their proximity sensors. This wall-following behavior is complemented with some use of the camera in order to avoid conflicts with other transporting robots and help localization in the arena. Robots use their camera when, based on the readings from their proximity sensors, there is the possibility that they might be entering or leaving the corridor, or when an obstacle is detected. Based on the colors detected in the captured images, the robot can distinguish between other robots and walls of different colors. By navigating in this manner through the arena, robots are able to pick up virtual payloads and deploy them.

5.1.2 Traffic Regulation

If the need for traffic regulation arises due to a conflict between two transporting robots in the corridor, two robots assume the institutional role of traffic regulators. The two traffic regulators place themselves at the opposite ends of the corridor so that each regulator can control the flow of transporting robots entering the corridor from one of the rooms (see Fig. 4-(b)). The goal of the regulators is to ensure that

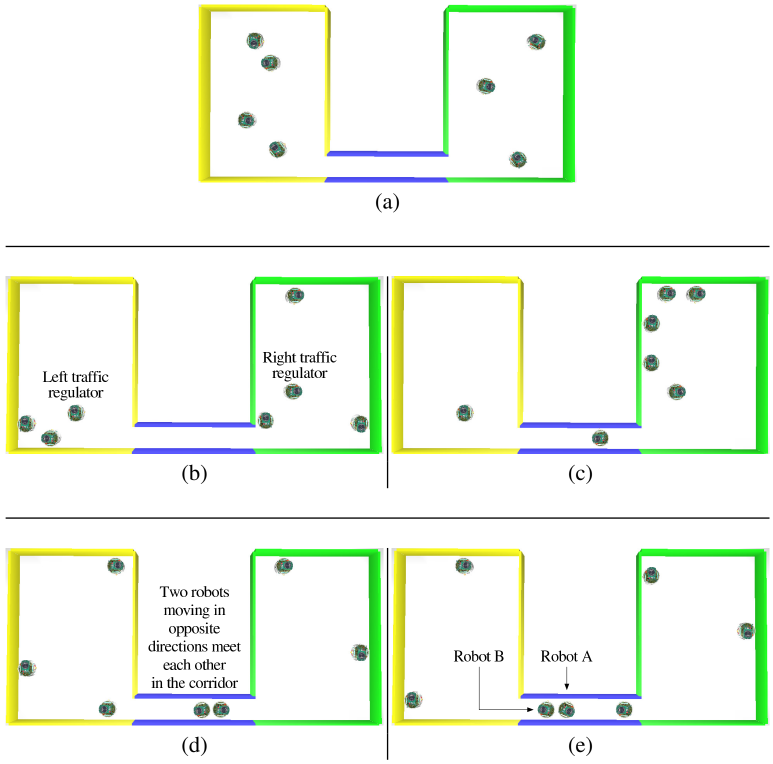

Fig. 4 Screenshots from submicroscopic simulations: (a) initial deployment of robots in the rooms; (b) regulators in their final positions at each entrance of the corridor; (c) queue formed behind right traffic regulator, while robot moves in the corridor. (d) two robots encounter each other in the corridor; (e) after adopting the role, robot A switches the role with robot B.

robots only move through the corridor in one direction at a time. The regulating robots are synchronized so that only one of them will let transporting robots enter the corridor from their respective rooms at any one time. The synchronization between the regulators is facilitated by an external program running on a *Webots* supervisor node, although it could also have been designed in a decentralized manner.

The regulators use their short range communication capabilities to emit messages to guide the transporting robots trying to enter the corridor. A traffic regulator periodically emits messages when it has to prevent transporting robots from entering the corridor from the room in which it is placed. Transporting robots have to be inside the short range communication radius of the regulators (set to 15 cm) to receive messages. If a transporting robot receives a message to stop, it will stop and begin to relay the stop message so other transporting robots behind it will stop too. As a result, the transporting robots will form a queue (see Fig. 4-(c)). When the first robot in the queue receives a message to proceed, it forwards the message to any robots that may be behind it, and the queued up robots will start to move.

5.1.3 Allocation of the Traffic Regulator Role

When two robots moving in opposite directions encounter one another in the corridor (see Fig. 4-(d)), they send a message to the supervisor to determine if they should adopt the role as traffic regulators. Each of the robots specifies from which room it came. If no other robot has yet assumed the role in the room specified by the robot, the supervisor instructs the robot to assume the traffic regulator role in that room. The robot, now that it has adopted the role, has to retreat to the room from which it came and place itself next to the entrance of the corridor. However, after two robots moving in opposite directions have assumed the role as regulators, other robots may already have entered the corridor and prevent them from navigating to the entrance of the corridor (see Fig. 4-(e)). In order to speed up conflict resolution in this case, the role is propagated to the last robot that entered the corridor from a given direction. Role propagation takes place in the following way: a traffic regulator (robot A) has been assigned the role, but has not yet navigated to the right location. During its retreat, robot A encounters another robot (robot B) in the corridor, both robots detect one another. Robot A stops, while robot B immediately sends a message to the supervisor in order to discover if it should assume the role as a traffic regulator. Despite the fact that a traffic regulator has already been assigned to the room from which robot B came (namely robot A), the regulator it is still located inside the corridor and not yet coordinating traffic. Thus, the supervisor sends a message to robot A to cancel the role assignment and instructs robot B to adopt the role instead. Robot A abandons the role, turns around and assumes the behavior of a regular transport robot.

After exiting the corridor, the regulator sends a message to the supervisor stating that it has made it outside of the corridor, preventing the supervisor from propagating the role further. The regulator navigates to a specific position at the entrance of corridor and sends another message to the supervisor stating that it is ready to regulate traffic. This specific position (see Fig. 4-(b)) is chosen to allow transporting robots to enter the corridor while at the same time being close enough to the regulator to receive the messages that it emits.

When the regulators in both rooms are ready, the regulation process begins. The supervisor sends messages to both regulators and instructs one of them to let transporting robots enter the corridor, while the other regulator is instructed to prevent robots from entering the corridor from its side. After a certain amount of time, the supervisor sends messages to both regulators to stop robots trying to enter from either side. This allows the corridor to clear before robots from the opposite direction are let through. After another fixed period of time, the supervisor sends a message to the regulators instructing them to allow traffic in the opposite direction of that from the last cycle. After a given number of these switches, both regulators abandon the role and the system goes back to the initial state. This is done so that all robots have a chance to accomplish the task and no robot has the role for the entire duration of the experiment.

5.2 Institutional Agent Controller

As before, when using the IAC methodology to design robotic controllers that implement our institutional approach our aim is to specify behaviors that have a social nature as institutions and behaviors that have an individual nature as the robots' individual behavior.

The individual behavior of the robots specifies how the task at hand is accomplished. Picking up virtual payloads and deploying them is a behavior that has an individual nature, since it relates the robot only to the environment in which it is located. A single robot could accomplish the deployment task, although performance would be critically reduced. Thus, we specify the individual behavior *Ind* as an EPN that accomplishes exactly that behavior.

The main social behavior of the corridor case study is the traffic regulator institutional role. This is clearly a behavior that has a social nature. We consider that this behavior is specified as an institution I_R that manages the role of traffic regulator. Its initial condition $initial_R$ is the detection of a conflict in the corridor and its final condition $final_R$ is the end of regulation (time limit). Since we do not want this behavior to be executed concurrently with any other behavior, the deontic operator of institution I_R will be $StopAll$. The EPN $Inst_R$ to be executed by the robots follows the sequence of actions described in Section 5.1.

However, the institutional role is not the only social behavior present. Institutional roles depend on other robots' behaviors, in the sense that other must recognize and/or permit such role playing by particular robots. A second social behavior present in the institutional approach to this task is the recognition and compliance with the traffic regulator. The behavior corresponds to an institution I_M that manages the reception of messages from the traffic regulators and their relay, which is implemented with the EPN $Inst_M$. Its initial condition $initial_M$ is the reception of a stop message and its final condition $final_M$ is the reception of a go message. We do not want this behavior to be executed concurrently with the individual behavior, so its deontic operator will be $StopInd$.

In Fig. 5 we show the higher layer composition of our two institutions and individual behavior. The IAC for this case study is the result of merging this net with the lower layer EPNs. More details, including the lower layer EPNs can be found in [17].

5.3 Self-Organized Approach

We implemented a different solution to our task which does not use institutional roles to regulate traffic. This solution is based on the principles of swarm robotics and the robots rely exclusively on self-organization to solve the task. Conflicts between robots moving in opposite directions in the corridor are solved in the following way: whenever a robot moving in one direction encounters a robot moving in the opposite direction in the corridor, it waits for a period of time proportional to

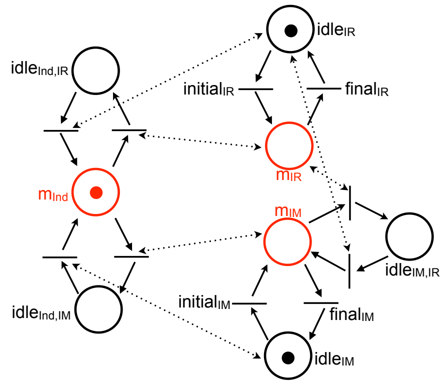

Fig. 5 Higher layer composition net of IAC for corridor case study. As before, dotted arcs represent bidirectional arcs. Places in red are macro places for behaviors in the lower layer. Place m_{Ind} represents the individual behavior Ind. Place m_{IR} represents institution I_R. Place m_{IM} represents institution I_M.

the time that it has been in the corridor. If, during this period, a waiting robot detects that the other robot gives up, turns around and moves back to the room it came from, the waiting robot continues to traverse the corridor. Otherwise, if the time proportional to the time the waiting robot has been in the corridor expires, the waiting robot turns around and heads back to the side of the arena from where it came. No further optimization of the self-organized approach was carried out. For instance, solutions using basic local communication could lead to better performances.

5.4 Experimental Setup

As with the previous case study, we use the *e-puck* robots as our robotic platform. We consider that each robot is endowed with two different forms of communication: short-range and long-range. Long-range communication can be achieved using Bluetooth, while short-range communication can be achieved using the *e-puck* proximity sensors as an infrared communication device.

We prepared different setups in order to evaluate how parameters such as the size of the robotic team and the length of the corridor affect the performance. Three different corridor lengths L (50 cm, 100 cm and 200 cm) were considered. For each corridor length, we ran experiments with different numbers of robots N (7 and 20 robots). For each of the six resulting setups, we performed 30 runs for both the institutional approach and for the self-organized approach. Each run had a duration of $T = 900$ seconds. Other parameters, such as the area of the rooms or the time intervals during with the regulators allow or stop robots entering the corridor, are directly dependent on N and L and are described in detail in [17].

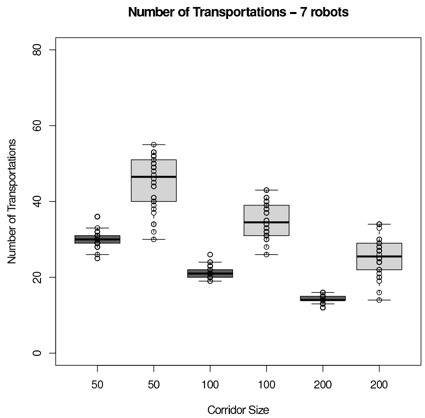

(a)

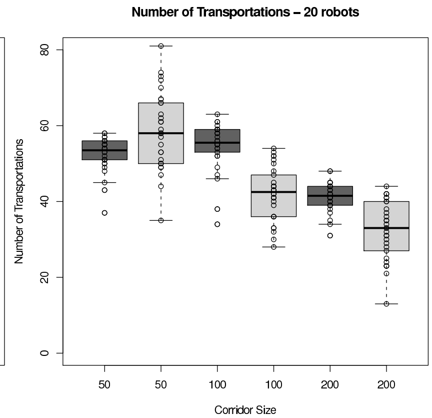

(b)

Fig. 6 Distribution of number of transportations for teams with (a) $N = 7$ robots and (b) $N = 20$ (institutional in dark grey, self-organized in light grey).

5.5 Results and Discussion

Our metric of interest is the *number of transportations*, described as the total number of successful virtual payload deployments (a pick up in the left room followed by deployment of the virtual payload in the right room) achieved by the team during a run of the experiment.

In Fig. 6, we display the distributions of the number of transportations for all six experimental setups. The results for the institutional approach are presented in dark grey while results for the self-organized approach are presented in light grey. The number of transportations decreases as the length of the corridor increases. This is naturally explained by the fact that the robots spend more time traversing the longer corridors.

When $N = 7$ (Fig. 6-(a)), we observe that the robots following the self-organized approach always manage to perform more transportations than the robots following the institutional approach. This is due to the fact that not all robots in the institutional approach are performing transportations. While in the self-organized approach all the robots are devoted to transporting virtual payload, in the institutional approach two of the robots instead assume the institutional role of traffic regulators. This means that some of the team's resources are spent on coordination. In small teams, a proportionally larger share of robots are dedicated to coordination (28.5% in the case of $N = 7$ robots with 2 traffic regulators). Moreover, in the self-organized approach conflicts are easily solved the first time a (proportionally) large group of robots meets in the corridor, resulting in the robots forming a line and thus making future conflicts rare. This emergent coordination allows the self-organized approach to perform more transportations when the team is small.

When $N = 20$ (Fig. 6-(b)), we observe an increase in the number of transportations for both approaches, although the increase for the institutional approach is considerably larger than for the self-organized approach. For $L = 50$ cm, there is not a significant difference between both approaches, except that the variance is much larger in the self-organized approach. However, for the longer corridors, the robots following the institutional approach perform more transportations. For teams with $N = 20$ robots, a smaller share of resources are dedicated to coordination (10% in this case).

Larger teams have a greater need for regulation than smaller ones, as they are more prone to conflicts occurring often, simply due to their larger number of robots. Since only two robots are devoted to the regulation at any time, larger teams spend less of their resources in coordination than smaller ones. Thus, larger teams have their need for regulation satisfied while allowing a larger share of robots to perform the transport task. The coordination of the team provided by the traffic regulators gives some advantage over the self-organized approach.

The larger variation, with respect to experiments with $N = 7$, in results for the institutional approach is due to the fact that, with a higher number of robots more conflicts occur in the corridor. It is less likely that the first two robots that encounter in the corridor eventually become regulators. Robots may switch the role between them multiple times, leading to a difference in the time that it takes before the traffic regulators effectively start coordinating the rest of the team (and therefore a difference in number of transportations).

For different sizes of the team and different corridor lengths, we observe that the variance of results is always smaller in the institutional approach than in the self-organized approach. This suggests that the regulation not only positively affects the performance of the system, but also its performance reliability. Other metrics such as the number of conflicts in the corridor and the time duration of transports were also studied. These results are described in [17].

6 Discussion on Social-Aware Coordination of Multi-Robot Systems in Human-Populated Environments

The corridor case study addresses three crucial issues for institutional roles: role allocation (how some robots start playing a role), role recognition (how robots recognize that some others are playing a role), and role permission (how robots permit other robots to play a role and behave accordingly). These three issues specify one of three elements involved in the ontology of institutional reality (according to [21, 22]): the assignment of status functions. Assigning status functions, that attribute some deontic powers, to objects or persons is one of the means human societies have to coordinate. In human societies, institutions exist because of the collective agreement of its members on the assignment of status functions. For instance, money, as a function, does not depend on the material chosen for banknotes or coins, but depends on the collective agreement to use such physical support to represent wealth and aid trade.

In the previous section we addressed the allocation, recognition, and permission, of the traffic regulator role. We saw that the allocation is made when a conflict occurs inside the corridor, and depends on whether other robots are already executing the role. The recognition of this role by transporting robots is made through a second institution, that associates the reception of a "stop" message with the knowledge that a regulator is in place and is coordinating the team. Role permission is also implied in this second institution, since it specifies how to behave accordingly. Every robot *can* play the role, so in fact every robot *has permission* to do so, meaning that transporting robots will accept the order of the regulator and conform to it. One can view this permission as the existence of a collective agreement by the robots to assign a status function to one particular robot playing the role. Institutional roles are one possible example of the type of coordination we discussed earlier: society-wide coordination.

Another relevant aspect of the corridor case study concerns the distinction between "role" and "individual." No robot is specifically designed to play the traffic regulator role. In the IR approach, playing a role is justified if there is a collective need for coordination, not as a right or an inherent feature of any individual. Particularly, if we consider scenarios taking place in human populated environments, institutional roles might also be played by humans, much in the same manner as they do when robots are not present. It is our intuition that, in certain scenarios, considering an IR approach will not only improve the coordination of the multi-robot system but will also facilitate the social interaction between robots and humans. Such intuition comes as a consequence of our goal to specify complex social interactions for multi-robot systems in a way that resembles the organizational aspects of human society.

The next main objective in the implementation and validation of the IR approach is to consider experiments in real-world scenarios populated with human actors. Robots will interact with these actors and such interaction should be based on coordination through common social rules, described as institutions. Such experiments will validate our intuition.

Let us consider a service robotics example scenario, focused on social robotics, using networked heterogeneous robots and sensors to interact with children, staff, and visitors, engaging in edutainment activities in the pediatric infirmary of a hospital. The infirmary environment includes not only bedrooms but also playing areas and also a school room. Besides being a realistic scenario, the ethical regulations enforced by hospitals for pediatric wards introduce important constraints, not only of a technical nature but also on what type of human-robot interactions are socially acceptable and how they are implemented. In order to deal with these constraints, we can consider institutional norms and institutional roles.

Institutional norms will allow the robots to comply with the social norms specified by both the ethics regulations and the staff of the pediatric infirmary. This will enforce that any coordination mechanism adopted by the robots is socially-aware. For a simple example of how such norms will be taken into account, consider a task where two robots must fetch a child from her bedroom and guide her to a playing area. This action will require coordination not only between the two robots but also between the robots and the child. Suppose, however, that another child in the same room is being observed by a staff member. In this case, an institutional norm can specify that the robots should not go into room (to avoid disturbing the medical examination) and should instead produce a visual sign to call the child to join them. Such norms can constrain the normal coordination procedure between robots and children to enforce social-awareness.

Institutional roles are also of great importance in this scenario. The robots must recognize that some individuals are playing a role that gives them some deontic powers not available to other individuals. However, there must also be recognition that these individuals might change but the role remains the same, for instance, different nurses have the same powers as long as they are on duty. In terms of human-robot interaction, recognizing that certain human actors have some deontic powers that others do not, not only makes the robots aware of the social structure of the interaction but also allows better coordination of the robots. For instance, a command to "go away" might originate different responses depending on who gives it. If issued by a doctor, a coordinated effort between all robots to move to a safe area might be in order, if issued by a child the robot should just leave the area nearby.

Due to the complexity of both the environment where they are set and the mission to be accomplished, social robotics scenarios in real-world environments populated with human actors can consider not only distributed control and coordination of a robotic system but also centralized behavioral planning. Social-aware coordination based on institutions can be present at both these levels. At the centralized level, institutions can be used to enforce the social-awareness of the plans developed. At the individual robot (distributed) level, institutions ensure that social norms are always taken into account in the execution of plans, even in the event of an unplanned change in the environment or a loss of communication with the centralized planner.

In the IAC methodology proposed for distributed control and coordination, the nature of behaviors (individual vs. social) and their implementation are extremely important. The implementations of behaviors as EPNs allows the combination of multiple institutions and individual behaviors into a single controller that executes a desired task while taking into account the institutional environment. The focus of centralized planning is on how to produce a sequence of behaviors for multiple robots that, when executed at appropriate times and locations, achieves a desired goal. From the perspective of the planner, the way in which behaviors are implemented (at the individual robot level) is not of importance but rather where, when and by whom those behaviors are executed. Thus, the EPN specification of institutions is not very relevant at the centralized planner level and the definition of institutions must be reformulated.

Given the different needs of the planner we consider the following definition for institutions (in the context of centralized planning) as a tuple (following previous work described in [24] that considered a similar definition):

$$\langle \text{activity}, \text{place}, \text{physical artefacts}, \text{roles}, \text{norms} \rangle \tag{1}$$

The determining factor in guiding the behavior of the robot is the *activity* in which humans want it to be involved. There must be a comprehensive and closed list of possible activities. In the case of the pediatric infirmary scenario, some examples are “good-morning room tour”, school-time, play-time.

Each activity can only be initiated at an appropriate *place*. A preparatory action may be necessary since the robot must be in a suitable place. For instance, was the robot at the corridor when accepting the command “it’s time for school”, it must first go to the school room. Certain activities can only take place at a specific location, while other activities can take place at different locations. Each activity must take place at a location which is appropriate from the humans’ point of view, so as not to be intrusive.

Then, a set of *physical artifacts* and a set of *roles* are used to confirm that the robot is at a suitable location for the intended activity, and that it has the human partners needed for that activity. For example, there is a whiteboard at school and nowhere else; there must be a teacher at school, although we can also find a teacher outside the school; certain activities can only take place in the presence of at least one child. *Norms* express the desired relations between robots and humans, using the institution as a coordination artifact.

Our research will proceed concurrently on three fronts: improving the IAC distributed methodology; introducing the IR approach in the context of centralized planning; and ensuring that both distributed and centralized approaches to IR are coherent in the design, representation and execution of institutions.

7 Conclusion

In this work, we describe and discuss several aspects of the IR approach to the coordination of multi-robot systems. Previously, we introduced a methodology based on the formalization of the central concept of IR - institutions - using EPNs, that allows us to design and execute robotic controllers (IAC) that produce behaviors capturing complex social interactions. Herein, we advance the IR approach in three fronts: we move forward in the validation of the IAC methodology by considering a real-world case study; we tackle a case study dealing with more complex social interactions and compare the IR approach with a self-organized approach; and we discuss how the IR approach might be applied in a particular real-world human-populated scenario.

First, we describe a real-world implementation of the wireless connected swarm case study, following an IR approach. We observed that such approach was able to maintain the wireless connectivity of a swarm of 40 real, resource-constrained robots. Results on connectivity show that, despite the high influence of noise in real world wireless communication, the overall swarm behavior implemented using an IR distributed control approach is able to maintain the expected connectivity.

With the corridor case study we have demonstrated how concepts from institutional robotics can be applied in a robotics task, focusing on one specific form of institution, namely the institutional role. We have shown that coordination artifacts set up as institutional roles can effectively help a robotic team organize and improve performance in a given task. Nevertheless, this is not true in all cases. For instance, we showed that for smaller teams, emergent coordination from a set of simple control rules is sufficient for the team to achieve a good performance. With the increase of the size of the robotic team, and the consequent decrease in proportion of robots devoted to institutional roles, we see benefits of using institutional roles, not only in the overall performance of the task but also in its reliability.

The next main objective in the implementation and validation of the IR approach is to consider experiments in real-world scenarios populated with human actors. This will be achieved in the scope of the MOnarCH[1] project. This project focuses on social robotics, dealing with the use of networked robots in an hospital setting, much the in same way as the example presented in Section 6. The development of the MOnarCH project will provide us with an ideal testbed to test our intuition that the IR approach will ease the effort of the transition to real-world environments populated with human actors and facilitate coordination with such actors in scenarios involving many-to-many human-robot interactions.

References

1. G. Beni. From Swarm Intelligence to Swarm Robotics. In E. Sahin and W. M. Spears, editors, *Swarm Robotics*, pages 1–9. Springer, 2005.
2. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.
3. C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Springer, second edition, 2008.
4. C. Cianci, X. Raemy, J. Pugh, and A. Martinoli. Communication in a swarm of miniature robots: The e-puck as an educational tool for swarm robotics. In *Simulation of Adaptive Behaviour (SAB-2006), Swarm Robotics Workshop*, pages 103–115, Rome, Italy, 2006.
5. R. Conte and C. Castelfranchi. *Cognitive and social action*. Londra: London University College of London Press, 1995.
6. N. Correll and A. Martinoli. Modeling and designing self-organized aggregation in a swarm of miniature robots. *International Journal of Robotics Research*, 30(5):615–626, 2011.
7. M. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-Based Multirobot Coordination: A Survey and Analysis. *Proceedings of the IEEE (Special Issue on Multirobot Coordination)*, 94(7):1257–1270, July 2006.
8. M. B. Dias, R. Zlot, M. Zinck, J. P. Gonzalez, and A. T. Stentz. A Versatile Implementation of the TraderBots Approach for Multirobot Coordination. In F. C. A. Groen, A. Bonarini, and N. Amato, editors, *Intelligent Autonomous Systems 8 (Proceedings of the International Conference on Intelligent Autonomous Systems IAS 2004)*, pages 325–334. IOS Press, 2004.
9. B. P. Gerkey and M. J. Mataric. Sold !: Auction Methods for Multirobot Coordination. *IEEE Transactions on Robotics*, 18(5):758–768, 2002.
10. G. M. Hodgson. What Is the Essence of Institutional Economics ? *Journal of Economic Issues*, 34(2):317–329, 2000.
11. G. M. Hodgson. What are Institutions? *Journal of Economic Issues*, XL(1):1–25, 2006.

[1] MOnarCH - Multi-Robot Cognitive Systems Operating in Hospitals, FP7-ICT European project. More info at `http://monarch-fp7.eu/`

12. P. Levis, S. Madden, J. Polastre, R. Szewczyk, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler. Tinyos: An operating system for sensor networks. In *Ambient Intelligence*. Springer Verlag, 2004.
13. T. Lochmatter, P. Roduit, C. Cianci, N. Correll, J. Jacot, and A. Martinoli. SwisTrack - a flexible open source tracking software for multi-agent systems. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4004–4010. Ieee, Sept. 2008.
14. O. Michel. Webots TM : Professional Mobile Robot Simulation. *Advanced Robotic*, 1(1):39–42, 2004.
15. F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J. Zufferey, D. Floreano, and A. Martinoli. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions*, volume 1, pages 59–65, 2009.
16. J. Nembrini, A. F. T. Winfield, and C. Melhuish. Minimalist Coherent Swarming of Wireless Networked Autonomous Mobile Robots. *From animals to animats 7: Proceedings of the seventh international conference on simulation of adaptive behavior*, pages 273–282, 2002.
17. J. N. Pereira. *Advancing Social Interactions Among Robots: an Institutional Economics-based Approach to Distributed Robotics Systems*. Phd thesis, IST-EPFL Joint Doctoral Initiative - Instituto Superior Técnico (IST), École Polytechnique Fédérale de Lausanne (EPFL), 2014.
18. J. N. Pereira, A. L. Christensen, P. Silva, and P. U. Lima. Coordination Through Institutional Roles in Robot Collectives (extended abstract). In S. e. van Der Hoek, Kaminka, Lespérance, Luck, editor, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems*, pages 1507–1508, Toronto, Canada, 2010.
19. J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli. An Experimental Study in Wireless Connectivity Maintenance Using up to 40 Robots Coordinated by an Institutional Robotics Approach. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5073–5079, 2013.
20. J. N. Pereira, P. Silva, P. U. Lima, and A. Martinoli. Formalization, Implementation, and Modeling of Institutional Controllers for Distributed Robotic Systems. *Artificial Life*, 20(1):127–141, 2014.
21. J. R. Searle. What is an institution? *Journal of Institutional Economics*, 1(1):1–22, June 2005.
22. J. R. Searle. Social Ontology : Some Basic Principles. *Anthropological Theory*, 6(March 2006):12–29, 2006.
23. P. Silva and P. U. Lima. Institutional Robotics. In *Proc. of ECAL 2007 - 9th European Conference on Artificial Life*, pages 157–164, Lisboa, Portugal, 2007.
24. P. Silva, R. Ventura, and P. U. Lima. Institutional environments. In *Proc. of Workshop AT2AI-6: From agent theory to agent implementation, AAMAS 2008 - 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 157–164, Estoril, Portugal, Dec. 2008.
25. L. Tummolini and C. Castelfranchi. The cognitive and behavioral mediation of institutions: Towards an account of institutional actions. *Cognitive Systems Research*, 7(2-3):307–323, June 2006.
26. D. Vail and M. Veloso. Dynamic Multi-Robot Coordination. In A. C. Schultz, L. E. Parker, and F. E. Schneider, editors, *Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II (Proceedings from the 2003 International Workshop on Multi-Robot Systems)*, pages 87–100. Springer, 2003.
27. A. Varga. Software tools for networking: OMNeT++. *IEEE Network Interactive*, 16(4), 2002.
28. D. Weyns, M. Schumacher, A. Ricci, M. Viroli, and T. Holvoet. Environment , a first-order abstraction in multiagent systems. *Knowledge Engineering Review*, 20(2):127–141, 2005.
29. A. F. T. Winfield, W. Liu, J. Nembrini, and A. Martinoli. Modelling a wireless connected swarm of mobile robots. *Swarm Intelligence*, 2(2-4):241–266, Sept. 2008.
30. A. F. T. Winfield and J. Nembrini. Safety in numbers : fault-tolerance in robot swarms. *International Journal of Modelling, Identification and Control*, 1(1):30–37, 2006.
31. L. Xu and A. T. Stentz. Market-based Coordination of Coupled Robot Systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2784–2789, 2011.