Representation and Execution of Social Plans through Human-Robot Collaboration

Lorenzo Nardi and Luca Iocchi

Dept. of Computer, Control and Management Engineering
Sapienza University of Rome, Italy

Abstract. The use of robots in people daily life and, accordingly, the requirement for a robot to behave in a socially acceptable way are getting more and more attention. However, although many progresses have been done in the last years, robots still have many limitations when they are required to share the environment with humans.

In this paper, we define the concept of *social plans* combining two main ideas: the definition of social behaviors for enabling a robot to live with humans and the establishment of a symbiotic relationship among robots and humans to overcome robots' limitations. Social plans are plans containing both robot and human actions and we provide an execution model for them where human actions are replaced by a human-robot collaboration scheme in which the robot actively drives the interaction with a human in order to obtain the desired effect.

A fully implemented system has been realized following this idea and different examples are provided in order to demonstrate the effectiveness of the approach.

1 Introduction

Recently, advancement of robotics encouraged a gradual move of robots from laboratories into people daily lives for acting as partners or assistants. In order to achieve this, robots should be able to share their working space with the people inhabiting it. The first requirement for a robot that wants to perform tasks closely to people is that it must be safe to humans; however, in the direction of being perceived by people as an actual partner rather than a solely mechanical tool, also human comfort and social acceptability should be considered.

In the last years, different robot social navigation systems have been developed to enable robots navigating into environments inhabited by people and interacting with them for accomplishing diverse and even complex tasks in a socially-acceptable way. These include systems for robots acting as an interactive museum tour-guide [1], escorting residents in nursing homes [6] and giving directions to the clients in malls [3]. However, most of them are task-specific and typically hard to generalize.

Further, current robotic systems still suffer from significant limitations at perception, reasoning and actuation due to factors like poor accuracy of sensors and actuators, high complexity and costs. These limitations cause several evident difficulties to develop complete and autonomous robotic systems able to perform even simple tasks, such as opening a door or bringing a cup of coffee.

M. Beetz et al. (Eds.): ICSR 2014, LNAI 8755, pp. 266–275, 2014.

In order to overcome robots' limitations, in this paper, we introduce a new general approach based on the idea of symbiotic robotics that allows robots for realizing tasks in a social way. In fact, as asserted by Rosenthal et al. [8], since many of robots' difficulties are easy tasks to humans with whom they share the working space, robots should be able to interact and collaborate with them for receiving the help needed. We design this collaboration through plans represented in Petri Net Plans (PNP) [11] that explicitly take into account human actions. We call these plans *social plans.*

The contributions of this article are in the definition and execution of social plans and in the explicit representation of human-robot collaboration. A full implementation of different services on an actual mobile robot has been realized to show the effectiveness of the presented approach.

2 Related Work

Over the last years, some general systems for robot social navigation have been implemented. They focus not only on paths and trajectories a robot should follow to reach its goal, but also on general norms for appearing as social entities to humans. In this regard, Pacchierotti et al. [7] developed a system to navigate in a hallway in which the Hall's concept of proxemic space was exploited. Kirby extended Pacchierotti's work realizing COMPANION framework [5], [4]: a more general system for realizing navigation tasks, able to take into account a wider variety of social cues. Here, a global optimal planner considers a set of constraints including some general social conventions to produce socially correct robot behaviors.

However, despite these systems allowing robots for a socially-acceptable navigation, current robots have significant limitations that affect the accomplishment of even simple tasks. In order to overcome this issue, Rosenthal et al. [8] introduced symbiotic relationships among robots and humans. Each individual involved in a symbiotic relationship performs distinct asynchronous actions and the results affect all individuals involved. Accordingly, when a human assists a robot, people will receive back the service provided from its task completion. The idea of symbiotic robotics has been further extended considering robots capable to ask for help to the actual occupants of the environment without any supervision, distributing the load of help to all the people living in the environment and obtaining assistance from the largest number of humans available in it [9].

We propose a new general approach to enable robots for accomplishing different social services based on the idea of collaboration among robots and humans who share the working space with them for overcoming their limitations. To this end, we introduce a robotic system that, resuming the key principles of symbiotic robotics and accounting a set of shared social conventions, enables to execute robot plans realizing socially-acceptable behaviours.

Cirillo et al. [2] assert that classical robot planning systems, in which the state of the world is only affected by the robot, are no longer applicable to robots sharing their working space with people. They define a human-aware robot task planning, in which robots consider the forecasted future human activities and adapt their action plans accordingly. Similarly, our approach to robot planning takes into account humans sharing the environment with robot; however, we define social plans in Petri Net Plan in which human actions are explicitly represented to realize social tasks. Then, human actions are transformed into sub-plans of robot operations realizing the collaboration among robots and people and allowing for the accomplishment of diverse and complex tasks.

3 Social Plans

In this article, we propose a new approach based on the idea of human-robot collaboration to enable a robot for accomplishing different social tasks. In particular, we are interested in the definition and execution of social plans and in the explicit representation of human-robot collaboration. In the following, we introduce the concept of social plans describing how they can be defined and executed using the PNP formalism.

3.1 Definition of Social Plans

In robot planning, given an action theory, an initial situation and a goal, a planner computes a plan that will be executed by the robot. However, if a robot is not able to perform some actions necessary to reach the goal, a plan may not exist. For example, it would not be possible for a robot with no arms to find a plan for reaching a goal that would require the opening of a door.

We define the concept of *social plan* as a plan that combines actions performed by robots and by humans. Including human actions, a social plan allows for an explicit representation of human-robot collaboration. Given the specification of robot and human actions, the initial situation and the goal, a social plan can be generated by any planner. Otherwise, these plans can be manually designed by the robot programmer using a suitable formalism. In this paper, we are not interested in how social plans can be generated, but we focus on their representation and execution.

Since a social plan includes actions that must be performed by humans who are not aware of the global plan of the robot, it cannot be directly executed by a robot. In order to address this issue, it is necessary to transform the social plan into a robot action plan. In this transformation, human actions are converted into subplans of robot operations which realize a human-robot collaboration scheme in which the robot actively asks for help to any human passing nearby.

3.2 Human-Robot PNP

In order to represent social plans, in this paper, we decided to use Petri Net Plan (PNP) formalism [11]. The motivations for this choice are: 1) the availability of definition and execution model of multi-agent PNP, that we can similarly apply to humans and robots [10]; 2) the higher expressiveness of the formalism with respect to other languages that allow for representing complex plans and complex forms of human-robot interaction; 3) the implementation[1] of the formalism in Robot Operating System (ROS) that allows for an easy development of actual robotic applications.

Given the high expressiveness of this formalism, no automatic planner is available to generate PNPs. However, the main concepts of social plans and of transformation into executable robot plans could be applied to other languages as well.

A social plan in PNP is defined similarly to a multi-robot PNP in which one of the agents acting in the global plan is a human. This plan is called Human-Robot PNP (*HR-PNP*). In a HR-PNP, human actions are represented in the same way as robot actions and are labeled as *H_X*, where *X* is the name of the action expected to be executed by a human. For example, *H_Open_door* denotes the action of opening a door by a human.

As discussed in the previous section, a HR-PNP cannot be directly executed by a robot because it includes some actions that have to be performed by a human. Therefore, a transformation of the social plan is needed for enabling a robot to actually execute the corresponding behavior. Since we defined a HR-PNP as a particular multi-robot PNP, this transformation will resume the scheme for transforming a multi-robot PNP into a single-robot PNP [10]. In fact, from a multi-robot PNP it is possible to automatically produce a set of single-robot PNPs by dividing the part of the plan relative to each robot. In the same way, we can transform a HR-PNP into an executable PNP by maintaining robot actions and converting those actions that should be executed by a human into PNP sub-plans composed of robot operations for interacting with humans, preserving the correctness of the whole plan.

Since we are interested in automatic transformation of HR-PNPs into executable PNPs, we have to consider a transformation method that can be generally valid for every human action. Accordingly, we realized a template in PNP composed of robot actions that can be applied for replacing each of the human actions included into a social plan. This transformation makes explicit the collaboration among humans and robot. In particular, the PNP template includes all those actions and conditions that allow a robot for establishing an interaction with people in a socially-acceptable way and receiving from them the help needed to achieve its goal.

The transformation scheme from a human action into a subplan of robot operations is illustrated in Figure 1. Here, it is possible to understand that the PNP template is based on three generic robot actions/subplans that can be implemented in different ways depending on the human action considered: preparation action (*PreA_X*), communication action (*CA_X*) and perception action (*PA_X*), where *X* is the name of the human action taken into account.

- *PreA_X* is the PNP action/subplan executed by the robot for preparing itself to receive the human help needed to perform its task. Here, the robot waits, detects and approaches humans to establish an interaction with them.
- *CA_X* is the PNP action/subplan that allows a robot to ask humans for the assistance needed. Here, a time-out is necessary to handle failures in the

[1] pnp.dis.uniroma1.it

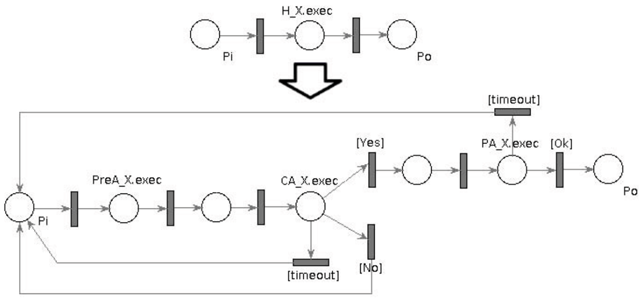

Fig. 1. Template PNP for transforming a human action in robot operations

communication. When a human refuse to help the robot or a failure is verified, the robot interrupts the interaction and tries to approach another person.
- *PA_X* is the PNP action/subplan that allows robot for sensing the space to ensure that the human accomplished the task he has been asked for. This action can be implemented exploiting diverse level of perceptions, from speech recognition to camera/laser data. Even in this case, a time-out is required to handle failures.

As illustrated in Figure 1, it is important to notice that the initial state (P_i) and the final state (P_o) of the original human action exactly coincide with the initial and final states of the template with which it is replaced. In fact, once the subplan defined in the template has been executed, the human action has been actually performed by some human who accepted to help the robot. In this way, human actions can be automatically transformed into subplans of robot operations without affecting the whole plan, given only the specifications of *PreA_X*, *CA_X* and *PA_X* that implement the human-robot collaboration for the action X.

Once all human actions of a HR-PNP have been transformed, an action plan that can be executed by a robot for realizing the corresponding behavior is obtained.

4 Implementation

Once the concept of social plan has been defined, we can introduce our system for realizing robot social behaviors describing its information workflow and providing an example of how does it work.

4.1 System Architecture

Currently, most of the systems demonstrating robot social tasks are specific for that particular task. Typically, these systems are effective but each of them is able to face just a few situations. Conversely, in this paper, we introduce a general

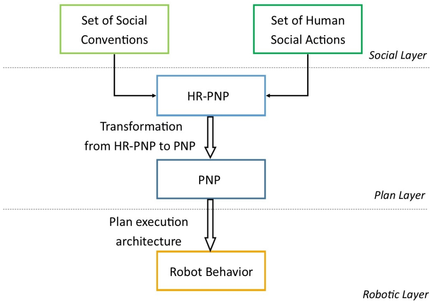

Fig. 2. Information workflow of our system

scheme that can be adopted and reproduced for easily realizing a large variety of robot social behaviors taking advantage from the collaboration with humans.

The information workflow of our robotic system is illustrated in Figure 2. It presents a layered architecture composed of three different levels.

- *Social Layer* includes a set of human actions, which defines the human actions that can be employed in designing social plans, and a set of social conventions, which includes different social norms shared among humans that range from navigation in public spaces to communication and allow a robot for interacting with humans in a socially-acceptable way.
- *Plan Layer* is the level in which a social plan for the desired service is designed in PNP considering the human actions defined in the social layer. Here, once a HR-PNP has been defined, the transformation into an executable PNP described in Section 3 takes place.
- *Robotic Layer* is the layer in which all different functionalities of the robot such as motion, speech and perception are combined together through ROS in order to actually execute the behavior corresponding to the PNP obtained from the Plan Layer.

4.2 An Example of Transformation of a HR-PNP

In order to provide an example of how our system works, we consider the simple HR-PNP illustrated in Figure 3a. Here, it is assumed to work with a mobile robot with no arms that, first, moves in front of a given closed door (*GoTo_door*); then, it is expected that a human opens the door (*H_Open_door*) in such a way the robot can pass through it (*Enter_door*).

Starting from this HR-PNP we want to obtain an equivalent executable PNP exploiting the method introduced in Section 3, i.e. replacing each human action that appears in the HR-PNP with a subplan of robot operations. To this end, we have just to transform *H_Open_door* as described before and to define the corresponding preparation, communication and perception actions that take into account the social conventions defined in the Social Layer of the system. A description of the definition of such actions is given in the following.

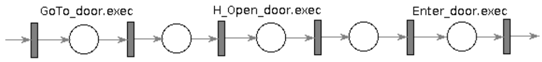

a) Open a door.

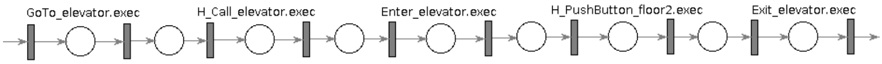

b) Take the elevator.

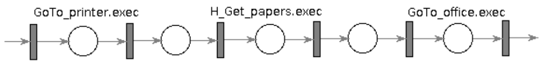

c) Get some papers.

Fig. 3. Examples of HR-PNPs for different robot behaviors

- *PreA_Open_door*: When the robot is in front of the closed door, since it is not able to open it by itself, it should look for help. For example, the robot can passively wait for a person who is passing nearby with whom trying to establish an interaction and to ask for assistance. Once a human has been detected by analysing its laser range finder readings, the robot approaches him greeting and turning in his direction.
- *CA_Open_door*: As a human has been approached, the robot tries to establish a conversation saying "Excuse me. I'm not able to open this door, can you help me?", which realizes the request for help by the robot. If the human refuses to give his help to the robot or the interaction fails for some reasons, it will look for another person who is available to help it.
- *PA_Open_door*: If the human accepts to help the robot, it senses the space through its laser range finder until the door has been opened. As the human opens the door, the robot thanks and greets him. Failure cases are handled by a timeout: if timeout is expired before the door has been opened, robot will restart the subplan asking for assistance to another person.

It is important to notice that we considered in social plans the concept of interrupt of PNP. When there is none in the surrounding of the robot available to help it, interrupt allows for suspending the current task and executing the next one, avoiding that the robot would be stuck forever waiting for assistance.

Table 1. Table for transforming some human actions into subplans of robot operations

Human Actions	**Preparation Actions**	**Communication Actions**	**Perception Actions**
Open a door.	Wait for a human and, when detected, approach him.	"I'm not able to open this door, can you help me?".	Check whether the door has been opened.
Call the elevator.	Wait for a human and, when detected, approach him.	"I'm not able to call the elevator, can you help me?".	Check whether the elevator door is open.
Push the elevator button.	None.	"I would like to go up to the second floor, can you press the button?".	Check whether the elevator door is open.
Bring papers from the printer.	Wait for a human and, when detected, approach him.	"Can you take papers from the printer and put them on my plate?".	Check that there are some papers on the plate.

In Table 1, preparation, communication and perception actions for transforming some human actions we considered in our experiments have been outlined. In particular, the first line reports the transformation we described above.

5 Use cases

We have implemented some use cases on an actual robot to demonstrate the effectiveness of our approach. The working space considered is our Department. Here, a mobile robot with no arms equipped with a laser range finder, a camera, a microphone and a speaker, who knows a priori the map of the environment and some semantic information on it, has performed different social tasks sharing the space with the humans working there. During the tests, some people without being informed a priori have been approached by the robot asking for their help. The videos of the robot accomplishing the diverse use cases are available at: `https://sites.google.com/site/robotsocialnavigation/videos`.

5.1 Navigation from the Hallway to the Auditorium

In the first use case, the robot should navigate from the hallway to the auditorium passing through two doors. Since the robot is not able to open a door, in order to face those situations in which it may find a closed one along its path, it is necessary to consider human intervention into the social plan (Figure 4a). This is exactly the situation we discussed in the previous section and the first video shows it. Here, a complete human-robot interaction is presented; however, if the human opens the door before the end of the interaction, the robot will be anyway able to take advantage of it.

5.2 Navigation from the First Floor to the Second Floor

In the second use case, the robot should reach the second floor from the first floor taking the elevator. To this end, the robot should call the lift and push the button corresponding to the second floor. However, the robot is not able to perform these two operations by its own and, therefore, the human intervention is needed. This is made explicit into the portion of social plan represented in Figure 3b that includes *H_CallElevator* and *H_PushButton_floor2* actions.

In the second video, it is shown how the action of calling the elevator has been replaced with the robot operations outlined in the second line of Table 1. Accordingly, when the robot reaches the elevator, it looks for someone passing there who can help it (*PreA_CallElevator*). As a human has been detected, the robot approaches him greeting and turning in his direction. Then, the human stops and the communication starts with the robot saying "Excuse me, I'm not able to call the elevator, can you help me?" (*CA_CallElevator*). When the human accepts and calls the elevator (Figure 4b), the robot thanks him and waits until the elevator door is open (*PA_CallElevator*). As the elevator arrives, it greets the human and gets in.

Although it is not shown in this video, also *H_PushButton_floor2* can be transformed into a subplan of robot actions (third line of Table 1) to enable the robot for reaching the second floor and completing its task.

a) Open a door. b) Call the elevator. c) Get some papers.

Fig. 4. Examples of tasks performed by a human for helping the robot

5.3 Bringing Papers from the Printer to Someone's Office

In the third use case, we foresee the possibility for a user to send a paper to be printed to the robot. Once the robot has received the paper, it chooses the printer and sends there the paper. As the paper has been printed, the robot takes it and brings it to the user's office. The interesting part of this task, corresponding to the plan in Figure 3c, is how the robot can collect the paper from the printer. In fact, since the robot is not able to grasp the paper, the human assistance (*H_Get_papers*) is necessary to achieve the goal.

In the third video, it is shown how *H_Get_papers* has been replaced by a subplan of robot actions (forth line of Table 1). Preparation action *PreA_Get_papers* is implemented by looking for and approaching a human. Then, the robot asks him: "Can you take papers from the printer and put them on my plate?" (*CA_Get_papers*). As the robot detects through the camera that the human have put the paper on its plate (*PA_Get_papers*) (Figure 4c), it thanks and greets him, and navigates towards the user's office. Once the robot arrives in front of the office door, it sends a Skype message to the user notifying that it is outside his door and he can came to take the paper. When the robot detects that there are no papers on its plate, it greets the user and takes charge of the next service.

6 Conclusions

In this paper we have presented the definition and execution of social plans as plans that combine robot and human actions and that are represented using PNP formalism. In order to be executed, human actions are transformed into PNP subplans explicitly representing a human-robot collaboration scheme driven by the robot. This formalism allows for describing complex tasks and complex human-robot interactions. The implementation of a set of use cases realized on a mobile robot in our Department has been used to demonstrate the effectiveness of the approach.

The proposed method shows that a robot behaving in a socially acceptable way can overcome its limitations by asking for help to humans. In contrast with previous forms of human-robot interaction where the human drives the conversations, in this approach, the robot is actively looking for human collaboration asking people for help to execute a task. We believe that if this request is reasonable with respect to the situation and it is made in a polite and socially acceptable way, people would not deny their help to the robot.

Several further studies should be considered along this line. First, the social norms that drive the robot behaviors may be explicitly represented and used to generate the social plans. At this moment these are considered by the human designer of the system. Second, extensive user studies should be done to evaluate the system when used by non-expert users (e.g., typical visitors of our offices). Third, more complex tasks and increased forms of social interactions that also take into account context may be explored. Although further research is needed in order to fully assess the proposed method, the reported activities show its feasibility and promising results.

References

1. Burgard, W., Cremers, A.B., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., Thrun, S.: Experiences with an interactive museum tour-guide robot. Artificial Intelligence 114(1), 3–55 (1999)
2. Cirillo, M., Karlsson, L., Saffiotti, A.: Human-aware task planning: an application to mobile robots. ACM Transactions on Intelligent Systems and Technology (TIST) 1(2), 15 (2010)
3. Gross, H.M., Boehme, H., Schröter, C., Mueller, S., Koenig, A., Einhorn, E., Martin, C., Merten, M., Bley, A.: Toomas: interactive shopping guide robots in everyday use-final implementation and experiences from long-term field trials. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009, pp. 2005–2012. IEEE (2009)
4. Kirby, R.: Social robot navigation. Ph.D. thesis, Carnegie Mellon University, The Robotics Institute (2010)
5. Kirby, R., Simmons, R., Forlizzi, J.: Companion: A constraint-optimizing method for person-acceptable navigation. In: The 18th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2009, pp. 607–612. IEEE (2009)
6. Montemerlo, M., Pineau, J., Roy, N., Thrun, S., Verma, V.: Experiences with a mobile robotic guide for the elderly. In: AAAI/IAAI, pp. 587–592 (2002)
7. Pacchierotti, E., Christensen, H.I., Jensfelt, P.: Embodied social interaction for service robots in hallway environments. In: Field and Service Robotics, pp. 293–304. Springer (2006)
8. Rosenthal, S., Biswas, J., Veloso, M.: An effective personal mobile robot agent through symbiotic human-robot interaction. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, vol. 1, pp. 915–922. International Foundation for Autonomous Agents and Multiagent Systems (2010)
9. Rosenthal, S., Veloso, M., Dey, A.K.: Is someone in this office available to help me? Journal of Intelligent & Robotic Systems 66(1-2), 205–221 (2012)
10. Ziparo, V., Iocchi, L., Lima, P., Nardi, D., Palamara, P.: Petri Net Plans - A framework for collaboration and coordination in multi-robot systems. Autonomous Agents and Multi-Agent Systems 23(3), 344–383 (2011)
11. Ziparo, V.A., Iocchi, L.: Petri net plans. In: Proceedings of Fourth International Workshop on Modelling of Objects, Components, and Agents (MOCA), pp. 267–290 (2006)