Behavior Coordination Mechanisms

– State-of-the-art

By

PAOLO PIRJANIAN

USC Robotics Research Laboratory
University of Southern California
Los Angels, CA 90089-0781
phone: +1 (213) 740-4523, fax: +1 (213) 740-7512
E-mail: paolo@robotics.usc.edu
URL: http://www-robotics.usc.edu/~paolo

October 7, 1999

Abstract

In behavior-based robotics the control of a robot is shared between a set of purposive perception-action units, called *behaviors*. Based on selective sensory information, each behavior produces immediate *reactions* to control the robot with respect to a particular *objective*, i.e., a narrow aspect of the robot's overall task such as obstacle avoidance or wall following. Behaviors with different and possibly incommensurable objectives may produce conflicting actions that are seemingly irreconcilable. Thus a major issue in the design of behavior-based control systems is the formulation of effective mechanisms for coordination of the behaviors' activities into strategies for rational and coherent behavior. This is known as the *action selection problem* (also refereed to as the behavior coordination problem) and is the primary focus of this overview paper. Numerous action selection mechanisms have been proposed over the last decade and the main objective of this document is to give a qualitative overview of these approaches.

1 Introduction

In behavior-based robotics the control of a robot is shared between a set of purposive perception-action units, called *behaviors*. Based on selective sensory information, each behavior produces immediate *reactions* to control the robot with respect to a particular *objective*, i.e., a narrow aspect of the robot's overall task such as obstacle avoidance or wall following. Behaviors with different and possibly incommensurable objectives may produce conflicting actions that are seemingly irreconcilable. Thus a major issue in the design of behavior-based control systems is the formulation of effective mechanisms for coordination of the behaviors' activities into strategies for rational and coherent behavior. This is known as the *action selection problem* (also refereed to as the behavior coordination problem) and is the primary focus of this overview paper. Numerous action selection mechanisms have been proposed over the last decade and the main objective of this document is to give a qualitative overview of these approaches.

The action selection problem has in the past been studied from several different points of view. The main research directions can be divided into the following fields, where action selection (or an equivalent problem) is a central issue:

- **Ethology**: From this point of view, the behavior of biological creatures has been investigated in an attempt to understand and ultimately model the action selection mechanisms used in biological systems.
- **Artificial Life**: Artificial Life ("AL" or "Alife") is the name given to the discipline that studies "natural" life by attempting to recreate biological phenomena from scratch within computers and other "artificial" media. Alife complements the traditional analytic approach of traditional biology with a synthetic approach in which, rather than studying biological phenomena by taking apart living organisms to see how they work, one attempts to put together systems that behave like living organisms.
- **Virtual reality & graphical simulation**: Coordination of the activities of avatars and artificial (autonomous) characters and their interaction with virtual environments and other real or virtual characters is an action selection problem. Such environments compromise a suitable testbed for studying action selection mechanisms.
- **Software agents**: Software agents are "intelligent" software components that in cooperation with other agents in a network (such as the Internet) retrieve information of interest to a given user. Action selection is central to the problem of how software agents communicate and cooperate.
- **Robotics (physical agents)**: From an engineering point of view we want to utilize robots to perform specific tasks. A major problem in the design and development of such robots is concerned with the action selection problem. The major focus of this thesis is on action selection in physical robots.

It is important to note that this list of research areas in action selection is not exhaustive, nor is it complete. Rather it lists, from a *behavior-based control* point of view, most relevant research areas that directly address the problem of action selection. Several other fields of research could arguably be added to the list. For example, *planning* in artificial intelligence is concerned with the problem of synthesizing a plan a course of actions that can lead to the accomplishment of some given task. However, planning systems are not behavior-based. Quite the contrary. Further, in planning systems once a plan has been synthesized the problem of action selection is trivial and therefore this discipline is not included in the list. Other fields of science that (can) contribute to solving the action selection problem include *decision making* and *operations research*. E.g., operations research is concerned with helping management with determining appropriate policies and actions in economic, environmental or other complex systems where tradeoffs between multiple objectives are necessary. These fields do not directly address the action selection problem, rather they address problems that are fundamentally very similar and can thus inspire new approaches to the problem of action selection. Later in this paper, we see examples of action selection mechanisms that are based on certain decision making methods, e.g., Bayesian decision making.

Coming back to the abovementioned list, many of the listed fields are interrelated and constitute a source of inspiration for one another. For example, computational action selection mechanisms in robotics are often inspired by biological and ethological models of action selection [10, 33]. Conversely, ethological models of action selection are by some researchers tested and validated by use of robots that simulate creatures [8]. However, the aim in robotics is different from that of, e.g., ethology. For instance emergent behaviors are interesting to study from an ethological point of view, which might not be the case for robotics from an engineering point of view. From an engineering point of view, emergent behaviors may not be desired because the aim is to construct robots that operate according to some given specifications and thus *unpredictable* emergent behaviors are unwanted. Thus emergent behaviors are interesting, from an engineering point of view, as long as they can be predicted.

Another major distinction between these different approaches to the action selection problem is concerned with *uncertainty*. In the physically embedded systems, the problem of uncertainty is much more problematic than it is for systems that are not embedded in the real world. E.g., virtual agents in virtual environments can usually assume perfect sensing and actuation which is not true for robots (physical agents). Thus action selection mechanisms that work in virtual environments might not work satisfactorily for agents interacting with the real physical world.

2 Outline of paper

In this paper we review some of the work done in addressing issues related to the action selection problem (ASP). But what is the action selection problem and why is it a hard problem? In section 3 the problem of action selection is described and through this description the major challenges of the ASP are pointed out. *Actions* are entities that are in focus when dealing with the ASP and thus in section 4 we attempt to describe what actions are. This allows us to classify action selection mechanisms with respect to the type of actions they consider. Using these findings we propose a taxonomy for action selection mechanisms in section 5. In the remainder of the paper we describe a representative set of existing action selection mechanisms and we conclude with a discussion.

In order to simplify the analysis of the systems and approaches presented later in this section we will emphasize their application to behavior-based mobile robot control. However, many of the presented approaches may have more general application than mobile robotics alone. Further, it is not intended to emphasize on any particular approach in itself, what is important is to highlight principles, ideas and techniques. In addition, for each approach a discussion of its main contribution to the action selection problem and an analysis of its main advantages and disadvantages is given. In section 13 the analysis of the approaches are used for a discussion of issues relevant to action selection.

3 The action selection problem

Maes gives the following definition of the Action Selection Problem (ASP)[33]:

> *"How can an agent select 'the most appropriate' or 'the most relevant' next action to take at a particular moment, when facing a particular situation?"*

However, due to constraints such as environmental complexity and unpredictability and due to the agent's limited resources (computational etc.) action selection cannot be completely 'rational' or optimal. Nonetheless, despite of sparse resources and unpredictabilities in its environment, we want the agent to operate 'appropriately' or select 'good enough' actions. Thus an immediate question that pops to our minds is *what is a 'good enough' action*? Simon's concept of *satisficing* endows the behavior of producing 'good enough' rather than optimal action for an agent with cognitive limitations. The term *bounded rationality* has been used to characterize a satisficing behavior by, e.g., [55]. (see also [62, 20]).

[33] lists the following requirements for an ASM that produces 'good enough' behavior:

- **Goal-orientedness**: it favors actions that contribute to one or several goals
- **Situatedness**: it favors actions that are relevant to the current situation
- **Persistence**: it favors actions that contribute to the ongoing goal
- **Planning**: it looks ahead to avoid hazardous situations
- **Robustness**: the system should degrade gracefully, i.e., the malfunction of one part should not cause the failure of the entire system.
- **Reactivity**: it provides fast/timely responses

These requirements define characteristics of 'good enough' behavior, however, they do not give an explicit answer to the question we pose. An explicit answer can enable us to define a qualitative measure of 'goodness' which is necessary to identify whether an action is good enough or not. An interesting ethological view of the action selection problem given in [69] provides an explicit qualitative answer to this question:

> *"This [ASP] is the overall problem for an animal of how to select its actions so as to maximize its future expected genetic fitness by surviving and reproducing, and by helping close genetic relatives to do likewise."*

While this ethological view of the action selection problem might be appropriate for some creatures, yet for others one might need to slightly reformulate the definition in order to generalize it to apply to more cases. Instead of having the objective to *maximize the expected genetic fitness* we can say that the agent should *maximizes its expected utility*. The term *utility* can then be defined to any function reflecting the desires and priorities of the agent. As you might notice this formulation corresponds to Bayesian Decision Theory [45], advocators of which believe that agents acting according to this theory display *rational* behavior. In the review of ASMs it is seen that only a few of the approaches define a qualitative measure for evaluating selected actions.

In [69] a set of ASMs are implemented on a simulator and their performances are compared. Based on this study Tyrrell adds a number of requirements to the ones listed by Maes.

- **Dealing with all types of sub-problem**: in Tyrrell's terminology *sub-problem* corresponds to abstraction level of the agent hierarchy described by Minsky (see section 4). This requirement states that it is necessary to have action selection at all levels of the hierarchy. Tyrrell thinks that the same ASM should handle all sub-problems.
- **Compromise actions**: the need to choose actions that are best for the collection of behaviors (or system objectives) rather than for individual behaviors.
- **Opportunism**: this in contrast to persistence should allow the robot to interrupt the ongoing goal and pursue a new one.

Some of these requirements are in conflict. For example, the requirement of persistence is in conflict with opportunism and planning is in conflict with reactivity. Therefore, an action selection mechanism should trade-off between these requirements.

4 Definition of *action*

The action selection problem is concerned with selecting an appropriate *action* among a set of possible ones. It is important to know what 'actions' are if a mechanism is to select among them. What is an action? Or more specifically, what types of actions are interesting in the context of robot control?

From Webster's Revised Unabridged Dictionary (1913)

> **Action**, n. [OF. action, L. actio, fr. agere to do. See Act.]
>
> 1. A process or condition of acting or moving, as opposed to rest; the doing of something; exertion of power or force, as when one body acts on another; the effect of power exerted on one body by another; agency; activity; operation; as, the action of heat; a man of action. ...
>
> 2. An act; a thing done; a deed; an enterprise. (pl.): Habitual deeds; hence, conduct; behavior; demeanor.
>
> The Lord is a Good of knowledge, and by him actions are weighed. –1 Sam. ii. 3.
>
> 3. The event or connected series of events, either real or imaginary, forming the subject of a play, poem, or other composition; the unfolding of the drama of events.
>
> 4. Movement; as, the horse has a spirited action.
>
> ...

According to this definition an *action* constitutes physical movements: "the condition of acting or moving, as opposed to rest". This kind of an action will for a robot correspond to motor movements such as steering a mobile robot in some heading direction. The above definition further include habitual deeds "hence, conduct; behavior; demeanor" which in behavior-based control would correspond to activation of a *behavior* best suited to the desires of the robot. When we speak of an action we also speak of an agent that performs the action. The desires of the robot (or an agent in general) are manifested through action. Schöpenhauer [60] accounts for the relationship between an agent's will, the external world and its actions. According to Schöpenhauer's philosophy, the existing world consists of *will* and *representation*. Will and representation are the most elementary aspects of the world and are linked through action. An agent's sensations are by its brain organized into percepts and thus a representation of the world. Action is a (perceivable) manifestation of the will under given circumstances. And the body is the instrument to furnish the will through actions in the perceivable world. Thus, *will* is understood as the quality that allows to choose between the possible actions.

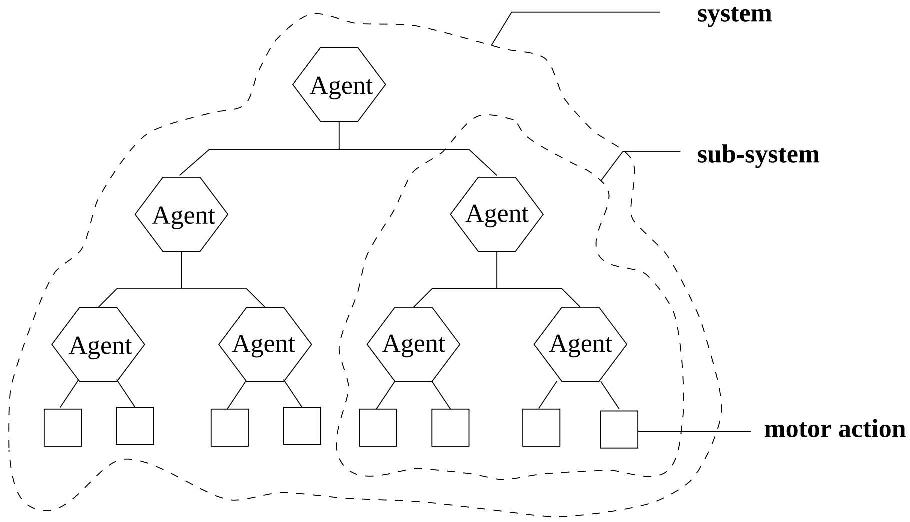

Figure 1: External view of agencies. The hierarchy of agents according to The Society of Mind. From an external view an agency can be divided into system, sub-system and motor actions.

But what is the relationship between actions causing movement and actions that correspond to the selection of behavioral processes (or simply behaviors)? In the Society of Mind, Minsky [37] draws the relation between actions and agents in a hierarchy (see figure 1). According to Minsky, an intelligent being consists of agents and agencies. Agents are organized in a hierarchy where abstract agents are build upon less abstract agents and so on. Each agent has a motive (or a task) which it pursues by activating and deactivating its subordinate agents. Looking at this hierarchy from outside it can be divided into agencies and sub-agencies. In [69] a *system*, *sub-system* terminology is used.

Within the agent, from one agent's point of view its subordinate agents correspond to the possible set of actions it has to select from. At the lowest level in the hierarchy there are motor actions which directly cause movement and interaction with the world.

Thus agents can have two roles in an agency: *actions* and *action selection mechanisms*. With respect to its subordinates an agent is an action selection mechanism and with respect to its superior an agent is viewed as an action. This is illustrated in figure 2.

5 A taxonomy for action selection mechanisms

Classification of existing action selection mechanisms (ASMs) into a number of logical groups can be useful. In the review of the literature we find two taxonomies ([32] and [56]) that suggest schemes for classification of ASMs.

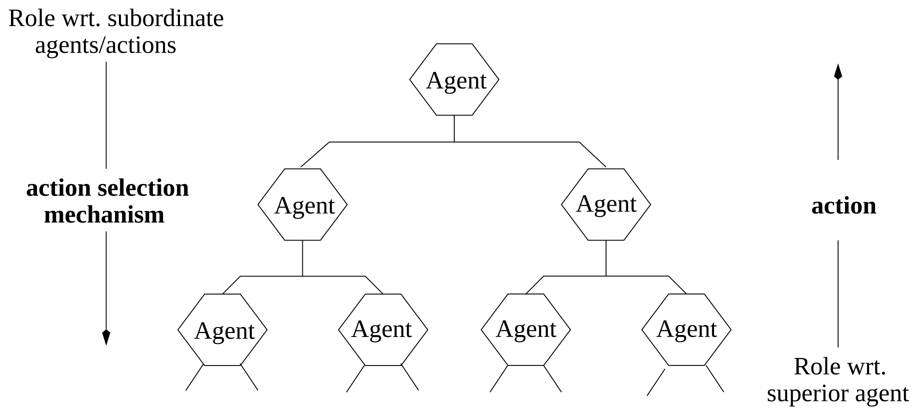

Figure 2: Internal view of agencies. The two roles of an agent in the agent hierarchy. From its own point of view an agent is an action selection mechanism whereas from a superior's point of view it is an action.

Figure 3 shows a classification of coordination mechanisms suggested in [32]. The classification divides coordination mechanisms into *state-based* and *continuous*. These terms are, however, not explained in [32], thus the following is our understanding of the terms. In a given state only a relevant subset of the robot's behavior repertoire needs to be activated. State-based coordination mechanisms select a set of behaviors that is adequately competent of handling the situation corresponding to the given state. This further leads to efficient use of sparse system resources. The authors of [32] further divide state-based coordination mechanisms into temporal sequencing and competitive.

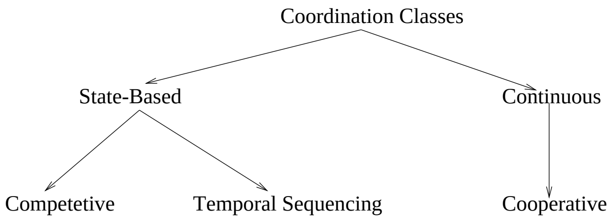

Figure 3: Classes of coordination mechanisms proposed in [32].

In [56] the ASMs are divided into *arbitration* and *command fusion* corresponding to MacKenzie's state-based and continuous approaches respectively. In Saffiotti's terminology each of these classes are concerned with the problems of *behavior coordination* and *command fusion* respectively. Behavior coordination is concerned with "how to decide which behavior to activate at each moment" and command fusion is concerned with "how to combine the results from different behaviors into one command to be sent to the robot's effectors" [56]. Note that this way of grouping ASMs corresponds to the roles of agents (as ASMs or motoric actions) discussed in the previous section.

Witnessed by these classifications it is evident that ASMs can be classified according to a number of relevant characteristics. However, its seems that one characteristic, namely the

ASM's ability to handle multiple behaviors simultaneously, is most discriminating and divides ASMs into two main groups: ASMs that can handle one behavior at a time and ASMs that can handle multiple behavioral constraints at a time. Following Saffiotti's taxonomy, we call these classes *arbitration* and *command fusion*, respectively. Arbitration ASMs, allow one or a set of behaviors at a time to take control for a period of time until another set of behaviors behavior is activated. Command fusion ASMs, allow multiple behaviors to contribute to the final control of the robot (agent). These ASMs are thus often suited for problems which are multiple objective in nature. For example, driving on a freeway and keeping the vehicle inside the lane, while avoiding collisions with surrounding vehicles is a multiple objective problem.

Various arbitration and command fusion mechanisms exist each or a group of which are implemented in various ways based on different mechanisms. Thus to further divide these two main classes we characterize them according to the mechanisms used for selecting the action. This is illustrated in figure 4.

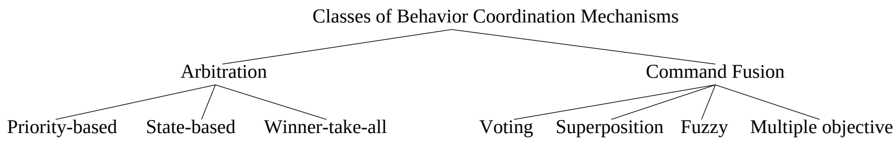

Figure 4: Classes of action selection mechanisms.

5.1 Arbitration ASMs

Arbitration mechanisms select one behavior, from a group of competing ones, and give it ultimate control of the system (the robot) until the next selection cycle. This approach is suitable for arbitrating between the set of active behaviors in accord with the system's changing objectives and requirements under varying conditions. It can focus use of scarce system resources (sensory, computational etc.) on tasks that are considered to be relevant.

Arbitration mechanisms for action selection can be divided into: *priority based*, *state-based* and *winner-take-all*. These are different techniques for selecting an appropriate action at any moment.

In *priority* based mechanisms an action is selected by a central module based on a priori assigned priorities. Thus behaviors with higher priorities are allowed to take control of the robot. *State-based* mechanisms select a set of behaviors that is adequately competent of handling the situation corresponding to the given state. Finally, in *winner-take-all* mechanisms action selection results from the interaction of a set of distributed behaviors that compete until one behavior wins the competition and takes control of the robot.

5.2 Command fusion ASMs

Command fusion mechanisms combine recommendations from multiple behaviors to form a control action that represents their consensus. Thus, this approach provides for a coordination scheme that allows all the behaviors to simultaneously contribute to the control of the system in a cooperative rather than a competitive manner.

Roughly speaking, command fusion techniques proceed in three steps:

1. **Action recommendations**: A module is implemented to generate recommended actions according to some behavioral criteria such as obstacle avoidance.
2. **Behavior aggregation**: The actions recommended by behaviors are combined according to some rule.
3. **Action selection**: An appropriate action is selected based on the combined recommendations.

Command fusion mechanisms can thus differ in the way they implement each of these steps. Steps 2 and 3 are often dependent on each other. In figure 4 the command fusion mechanisms are divided according to the way steps 2 and 3 are performed.

Voting techniques interpret the output of each behavior as votes, which are combined by tallying the votes, and then selecting the action which has received the maximum number of votes. *Fuzzy* command fusion mechanisms are very similar to voting techniques, however, fuzzy inferencing techniques are used for implementing steps 1 through 3. *Superposition* techniques combine behavior recommendations using linear combinations. Finally, there is the class of *multiple objective behavior coordination* [1], which provide a formal decision-theoretic approach to making decisions based on multiple objective decision theory.

Not included in this taxonomy are the class of *learning* techniques for action selection. These techniques can be used to 'learn' how to select an action. The most popular technique in this category is reinforcement learning.

6 Priority-based arbitration

6.1 The subsumption architecture [Brooks 86]

The subsumption architecture (SA) [10] promoted by Rodney Brooks, consists of a series of behaviors, which constitute a network of hardwired finite state machines. Action selection consists of higher-level behaviors overriding (or subsuming) the output of lower-level behaviors.

In the subsumption architecture Brooks defines *levels of competence* which is an informal specification of a desired class of behavior at different abstraction levels:

1. Avoid contact with objects
2. Wander aimlessly around
3. Explore the world
4. ...

A higher level of competence implies a more specific desired class of behaviors. Brooks, argues that adhering to the principles of the subsumption architecture one can incrementally build a system. A complete robot control system is built, which achieves one level of competence, e.g., avoid contact with objects. Once the designer is satisfied with the performance of this level he/she builds the 2nd level of competence. Each level is able to examine data from the levels below and is also permitted to inject data into the internal interfaces of the levels below thus suppressing normal data flow.

[1] Also known as multiple objective action selection mechanisms

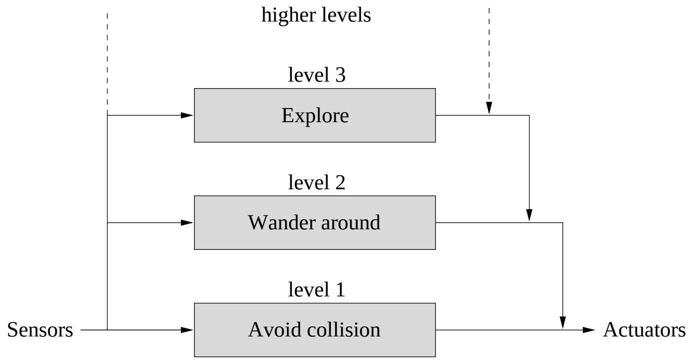

Figure 5: A subsumption architecture with a layer of levels of competence for a mobile robot.

Each level of competence in the subsumption architecture is implemented using a finite state machine augmented with timing elements and registers (denoted AFSMs, augmented finite state machines). Behaviors are usually constructed by the Behavior Language [11] which is a high level language for writing subsumption programs.

In the subsumption architecture (SA), modules/behaviors cannot cooperate - each layer is built on its own "hardware" totally from scratch. The control system is hard-wired directly in the structure of the behaviors and their interconnections, and can thus not be altered without redesigning the system.

It can be shown that a combinatorial explosion in the number of possible behaviors can occur when composing behaviors according to the subsumption architecture [3, 65]. However, in [15] it is demonstrated how the subsumption approach can be used to construct controllers for complex systems. The SA controller is designed on a single 15MHz CPU for a 19 DOF insect-like robot with 60 sensors. Another important limitation of SA is the lack of explicit goals and its goal-handling capabilities. The functionality of a SA based robot is viewed as an *emergent* property of the interaction of its behaviors and depends on the static and dynamic properties of the environment. In certain cases it would be convenient to have more direct control of the robot, however, typically there is no way to externally guide such robots.

The subsumption architecture has contributed to a fruitful debate in many fields including artificial intelligence and robotics. As a result of this debate, many ideas related to reactive and behavior based systems were investigated by many researchers, where its advantages and disadvantages were explored and demonstrated.

7 State-Based arbitration

7.1 Discrete Event Systems [Koŝecká 1993]

In this framework the interaction of behaviors are formalized by the DES formalism in a modular and hierarchical fashion [23]. Behavior selection is done using state-transition, where upon detection of a certain event a shift is made to a new state and thus a new behavior. Using this formalism, systems are modeled in terms of *finite state automata* (FSA), where *states* correspond to execution of actions/behaviors and *events*, which correspond to *observations* and *actions*, cause transitions between the states. See example FSA in figure 6.

The agent and its interaction with the environment are modeled by a FSA and is denoted the *plant*. The design objective is to specify an additional FSA, called the *supervisor*, that interacts with and modifies the open-loop behavior of the plant in order to assure correct behavior of the plant in accordance with some specified constraints. The supervisor achieves its objective by enabling and disabling the events which are *controllable* to modify the state transition function of the plant so that only desired subsets of event sequences are generated by the plant model.

In practice each actuator and perceptual component of the physical agent is modeled by a FSA called the *fundamental model*. Then to compose more complex behaviors from the fundamental models, a supervisor is synthesized that controls their interaction with each other and the environment. For complex systems and tasks this modeling task can, however, be very tedious and complex. Thus Christensen and Košecká have developed linguistic tools that automatically can transfer textual descriptions of plans to models and FSAs that implement the desired behavior [24]. This approach, denoted the *task specification language approach*, uses the *robot schema* ($\mathcal{RS}$) model of Lyons [31], which is a general process algebra. The transformation from task description to FSA models is performed in three steps. First a FSA is defined for each operator of the $\mathcal{RS}$ model. Then the FSAs are combined to build an FSA modeling the specified task, which finally is used to synthesize a supervisor for that task.

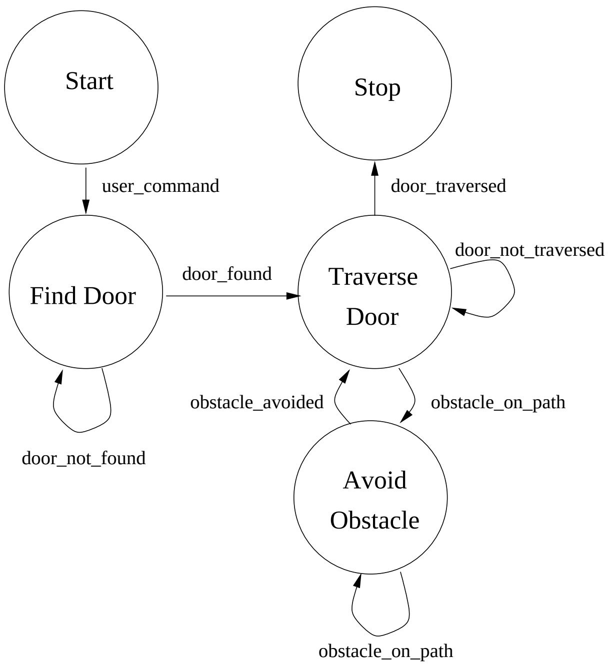

Figure 6: An example FSA encoding a door traversal behavior.

7.2 Temporal Sequencing [Arkin & MacKenzie 94]

The temporal sequencing approach also known as perceptual sequencing and sequenced coordination [6] is very similar to the discrete event system approach. A finite-state automaton is used for formulating a mechanism for sequencing between a series of behaviors based on perceptual triggers. At each state a distinct behavior is activated and perceptual triggers cause transitions from one state to another. The knowledge of when and how to activate which behavior is implemented in the behavior and hardware definition language BHDL. In this encoding scheme, states correspond to motor schemas (see section 12.2 for description of schemas) (e.g., avoid obstacle) and state transitions, also called perceptual triggers, correspond to perceptual schemas (e.g., detect obstacle) instances of which run in parallel with other schemas and signal a schema manager of state changes. Figure 6 depicts a simple FSA that consists of 5 states: start, find door, traverse door, avoid obstacle and stop. A user command triggers the robot to start looking for a door, when the door is found it starts traversing the door. During this state if an obstacle is detected on the robot path an obstacle avoidance behavior is invoked until obstacle clearance is achieved. The robot stops when the door has been successfully traversed.

This behavior coordination technique together with the motor schemas approach (see section 12.2) has been used in design of a series of robotic demonstrators including cooperating robots [7].

7.3 Bayesian Decision Analysis [Kristensen 1996]

The approach presented in [26] addresses the related problem of *sensor selection*, i.e., which sensors to use for which purpose. Sensor selection can be considered a special case of action selection, where the actions are certain sensor operations. It operates according to the purposive paradigm, i.e., the system consists of a set of purposive modules similar to behaviors. The task that is considered, (see figure 7) is to decide which sensors to allocate to which purposive modules in order to accomplish a given task which is declared by the *mission planning* module.

The approach is based on Bayesian decision theory and utility theory. The objective is to choose the action that *maximizes the expected utility* of the agent. In other words, sensing actions are evaluated in a cost/benefit manner where each sensing action is associated with a certain cost or expense and benefit is associated with the information provided by the sensing action. This approach has the view that planning ahead of time in dynamic environments makes little sense thus myopic planning is used, where only one step ahead of time is planned.

The structure of the decision problem is depicted in figure 8. The root node represents the decision task of sensor selection. The $m + 1$ arcs connected to the root denote the situation where the sensor is allocated to one of m purposive modules or not allocated to any module at all. The latter corresponds to the case where the cost of sensing would be higher than the expected payoff. The *report* node represents the random outcome, x_j, of the sensing action. To each report node, is connected an *actuator action* node which represents the final consequence of action that will lead to the completion of the task. The *state* node represents the state, z_l, that the world will assume after actuator action a_k. With each action-state pair is associated a utility, $U(a_k, z_l)$ that indicates the payoff from selecting action a_k when the state of affairs is given by z_l. For example, going through the doorway when the door is closed should have a low utility and going through the doorway when the door is open should have a high utility.

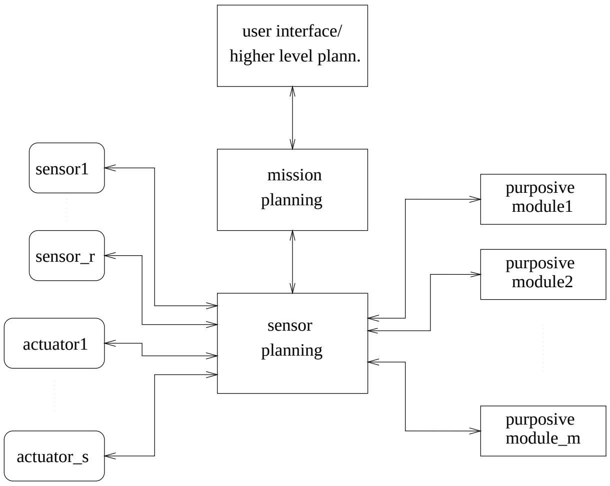

Figure 7: The Architecture used in the sensor selection approach in [26]. Certain sensors, e.g., active sensors, can be accessed by one purposive module at a time thus the *sensor planning* module has to resolve sensor contentions.

Thus at each cycle the planner decides which sensor to grant to which requesting purposive module, which in turn gathers information by sensing. This control behavior is depicted in figure 9 using a finite state automaton. When it no longer pays off to collect information, i.e., when A_0 is optimal, the planner selects an actuator action to perform, which consequently reveals what utility is earned.

In practice this sensor planning framework should have access to the following components:

- Probabilistic model of the purposive modules. This model is specified in terms of conditional probabilities that model how likely it is for a module to report the correct state of the environment.
- Probabilities, that express a priori information about the reports and the state of the world.
- Utilities and cost that reflect what it is worth performing an action in a given world state and the cost of doing so.

The construction of the probabilistic model of the purposive modules is usually done experimentally through a number of experiments. A set of experiments should be conducted for each purposive module in the target environment. Significant changes in the environment usually mean that the model should be regenerated. The utilities and costs of actions are usually determined heuristically and through fine-tuning until the desired behavior is reflected. The utilities in some sense represent preferences on the set of actions.

This approach has been applied to an indoor mobile robot navigation system, where the robot autonomously navigates from one room to another in an office like building avoiding obstacles on its way.

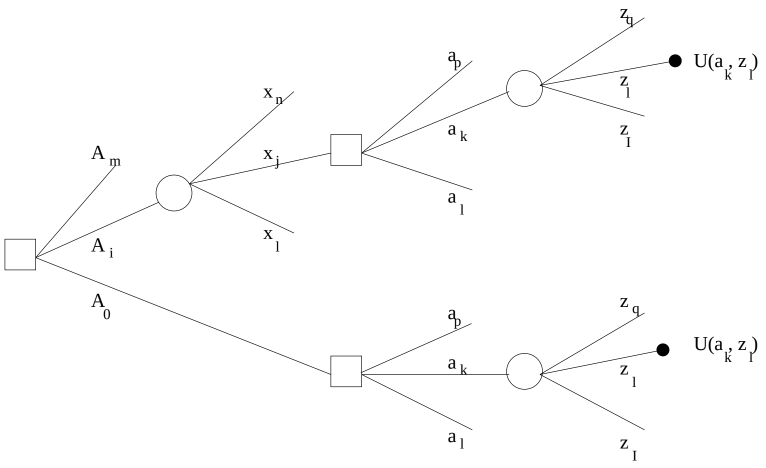

Figure 8: The standard structure of a decision tree, adapted from [45].

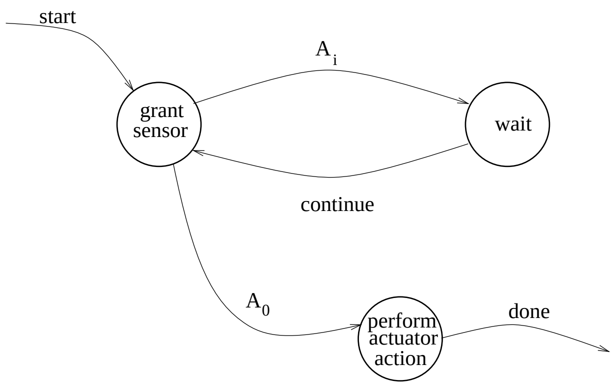

Figure 9: The control structure for a planning task, adapted from [27]. A_i means all actions except from A_0. While in the wait state the planner waits for sensor requests.

7.4 Reinforcement Learning approaches to action selection

A fundamentally different approach to action selection is to learn the action selection mechanism. Several approaches to learning of action selection exist and the most promising of the techniques is reinforcement learning. Basically reinforcement learning is a way to induce, based on trial and error, a perception to action mapping that maximizes some reward. The robot learns the perception-action mapping, known as a policy, by exploring actions that subsequently lead to some reward. The reward function is designed so to encourage desired behaviors and suppress unwanted actions. Thus the robot will select actions that maximize the expected reward. The design of the reward function can thus be as difficult, as hand coding of behaviors, for complex problems. Further, as the complexity of the problems scales up, the size of the state-space increases dramatically thus making the mechanisms practically intractable.

7.4.1 Hierarchical Q-learning [Lin 93]

To avoid the scaling problem of reinforcement learning for complex problems Lin proposes an approach that breaks down a complex task into a set of simpler tasks [30]. Much like the behavior-based approach to robot control, the learning problem is divided into a set of simpler problems each learned separately by a Q-learning module. Further, a Q-learning action selection mechanism is defined for arbitrating between the existing n modules. I.e., upon observing a state, x, the arbitration module selects a module, A_i, which then executes an action.

7.4.2 W-learning [Humphrys 97]

W-learning is a somewhat different reinforcement learning approach for action selection presented in [19]. In this approach each module/behavior recommends an action with some **W**eight (hence the term W-learning) and then the action with the highest weight is selected and executed. The W-values of each behavior are then modified based on the difference between the winning action and the action desired by the behavior. [19] proposes several ways of updating the W-values based on: minimize the worst unhappiness, maximize the best happiness etc. The results of this approach are demonstrated in a simulated environment of a house keeping robot that has to attend to several tasks including putting off fire, avoiding a predator etc.

8 Winner-take-all arbitration

8.1 Activation Networks [Maes 1989]

In this approach the system consists of a set of behaviors or competence modules which are connected to form a network. This approach is based on a society-of-mind type of idea [37], and it actually implements a difference-engine, where a set of behaviors reduce the difference between the system's present state and a goal state. This technique is called the Means-Ends Analysis and was first introduced in the General Problem Solver [40]. In the network each behavior is represented by a tuple $(c_i, a_i, d_i, \alpha_i)$ describing: *1)* the preconditions under which it is executable (i.e., can be applied), *2)* the effects after successful execution in form of an add-list a_i and delete-list d_i and *3)* activation level, α_i, which is a measure of applicability of the behavior. When the activation level of an executable behavior exceeds a specified threshold, it is selected to furnish its action. Evidentially this action selection mechanisms deals only with selection of behaviors and not with motor actions. When a behavior is selected it will perform the most appropriate action from its point of view. The activation level of the behaviors is influenced by an external injection and removal of activation energy from the network as well as internal exchange of activation energy between the modules within the network.

External sources of activation:

- Activation by the state: Activation energy is spread from the state of the environment to behaviors that match the current state[2]. ϕ is the amount of activation energy injected into the network for a valid proposition of the current state. To allow modules that

[2] A behavior is said to match the current state if at least one of its preconditions are true

use the same precondition to compete the amount of activation energy, ϕ, is divided equally between them.

- Activation by the goals: Each goal injects activation to the behaviors that as a consequence of their execution can fulfill the goals, i.e., one of the goals belongs to the behaviors add-list. γ is a parameter that determines the amount of activation energy injected into the network by a goal and is divided between behaviors that can achieve the goal.
- Inhibition by protected goals: The activation level of a behavior is decreased if it can undo a goal that has already been achieved. An already achieved goal (a protected goal) can be undone by a behavior that has the goal in its delete-list. δ is the amount of activation energy removed from the network by a protected goal.

Internal sources of activation:

- Activation of successors: An executable behavior, b, (i.e., one with satisfied preconditions) spreads activation to behaviors with preconditions that will become true after activation of b. For each true proposition in its add-list a behavior increases its successors' activation level by $\alpha\frac{\phi}{\gamma}$.
- Activation of predecessors: A behavior b that is not executable spreads activation to behaviors that may make it executable, i.e. behaviors which when executed will make some of b's preconditions true. For each false precondition a behavior increases its predecessors' activation level by α.
- Inhibition of conflictors: Every behavior decreases the activation level of behaviors, called conflictors, that can make its preconditions false. For each true precondition a behavior removes $\alpha\frac{\delta}{\gamma}$ from its conflictors.The amount of internal activation transfer to a behavior is divided by n, which is the number of propositions in the precondition list, add-list or delete-list depending on the type of activation energy exchange.

The activation energy exchange takes place continuously and at each cycle an executable behavior with the highest activation level is selected. Figure 10 illustrates an example activation network taken from [34]. If this behavior's activation level is above some threshold then it is allowed to execute otherwise the threshold is decreased by 10% and the cycle is repeated until some behavior is selected.

By exchange of activation energy the behaviors compete and cooperate to select an action (a behavior that performs an action) that is most appropriate to the current environment state and system goals. The way activation energy is exchanged between the behaviors, favors selection of a sequence of behaviors that can transform the current state of the environment to a goal state. This sequence is, however, not determined once and for all and then followed as in traditional AI planners. Unlike traditional planners, this system does not have an explicit representation of plans, neither does it implement a specific search algorithm to construct plans. Rather the system emergently chooses and performs the next step of the sequence, which is influenced not only by the system goals but also the state of the environment at each step. The selected actions correspond to system subgoals which bring the system closer to one of its goals. This does, not imply that the system will not fluctuate between different goals. The system is designed to be situation driven as well as goal oriented, so the system will be opportunistic and choose the shortest sequence of actions that can bring the system in a goal state. The system's sensitivity towards goals and the current state of the environment

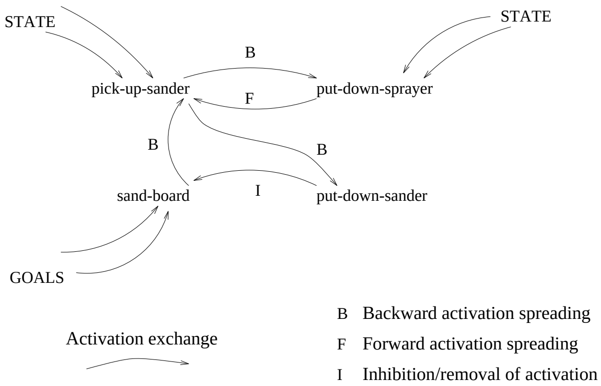

Figure 10: Example of an activation network. Adapted from [34].

can be influenced by the ratio of γ to ϕ. Recall that γ is the amount of activation energy injected into the network by a goal and ϕ is the amount of activation energy injected by state propositions. Thus an increase in γ makes the system more goal oriented and vice-versa, further an increase in ϕ makes the system more sensitive to the current state and vice-versa. There are no clear guidelines for choosing these parameters, they depend on the problem, specific activation network and its size etc.

This particular approach has been applied to simulated problems and has not been evaluated on real-world scenarios, which could help discover its strengths and limitations. One shortcoming of this approach might be that the state of the environment in systems do not include any uncertainty thus the system assumes perfect sensors etc. However, due to its distributed nature the system is fault-tolerant since the failure of a single module/behavior does not necessarily lead to the failure of the entire system.

9 Voting based command fusion

9.1 DAMN [Rosenblatt & Payton 89-96]

DAMN is a Distributed Architecture for Mobile Navigation [53], which consists of a set of behaviors, that pursue the system goals based on the current state of the environment. Each behavior votes for or against the set of actions constituting the possible set of actions of the agent. Behavior votes represent each behavior's preference over the set of possible actions. Then a 'voter' selects the 'best' action, which is chosen as the action with the maximum weighted sum of the received votes. The architecture is sketched in figure 11. Each behavior is assigned a weight by the mode manager. These weights reflect the relative importance or priority of the behavior in a given context. Rosenblatt, however, does not describe how the mode manager determines the weights.

Note that a behavior can be a reactive behavior as well as a planning module. In this way not only one behavior but several behaviors control the course of actions so that the system can simultaneously attend to several goals/subgoals and constraints of multiple independent behaviors. The *preference combination* algorithm implemented by the 'voter' is ad-hoc and the inventors of DAMN realize that there is a need for a voting scheme with well defined semantics and a predictable behavior. Being highly distributed, DAMN is tolerant to many failures, however, the 'voter' introduces a single point of failure.

The mechanism(s) for deliberation and purposive goal achievement in DAMN operates no different than all other components of the system. Like other system components the deliberative components declare their preferences by voting for or against the set of possible actions. In this way the preferences of the deliberative modules bias, rather than dictate, the course of actions towards a plan that can lead to a goal state. Thus the system is not restricted by any hierarchical constraints, where homogeneous modules are layered in the same layer. Rather all modules, homogeneous as well as heterogeneous, communicate their intentions using a common interface. This architecture takes the position, introduced by [1], that plans should be viewed as resources like any other resource, denoted the *plan-as-resource* view, rather than being interpreted as strict programs that should be followed word to word, denoted as *plan-as-program*.

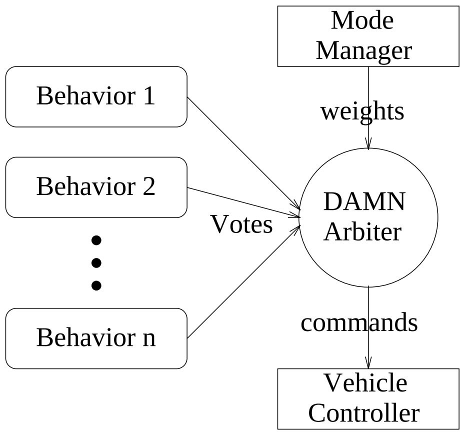

Figure 11: The Distributed Architecture for Mobile Navigation. Adapted from [53].

The advantages of preferential behaviors are demonstrated in [44], where it is argued that preferential outputs decrease the loss of information and serve a suitable interface between behaviors regardless of their level of planning. Further in [43] the authors demonstrate that this level of abstraction is well suited for integration of high-level deliberative planners with low-level reactive behaviors.

DAMN has been successfully applied to in-door as well as out-door real-world applications such as the out-door CMU Navlab vehicles, and indoor real and simulated robots at the Hughes Research Labs. In [53] the DAMN architecture is used for control of the heading and speed control variables of a mobile robot and in [54] the technique is used for choosing the filed of view for a motorized camera head. Further, [69] investigates and compares a number of ASMs, including DAMN and Maes' Activation Networks, in a simulated environment. The experimental results show that DAMN is superior to the other ASMs it is compared to.

9.2 SAMBA - Action maps approach [Riekki & Röning 96]

The action selection mechanism proposed in [52] is largely influenced by the ideas put forward by Payton and Rosenblatt in conjunction with the DAMN architecture. The work in [52] is seen as an elaboration of and an attempt to systematize the work done by Payton and Rosenblatt.

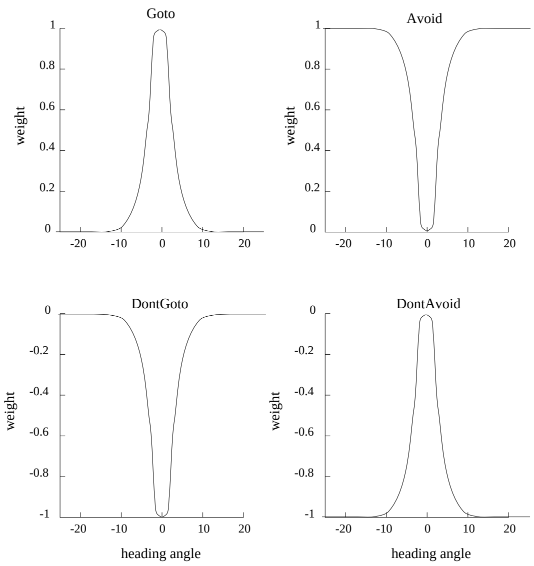

Figure 12: Two-dimensional primitive action map types.

In this work, primitive behaviors produce reactions in form of primitive action maps. An action map specifies preferences for each possible action, much like in the DAMN architecture. However, behavior outputs are generated from four primitive action map types as seen in figure 12. In figure 12 preferences are described for each heading angle, however, in general action maps can be defined on multi-dimensional control parameters. The action maps generated by the behaviors are forwarded to a command arbiter that combines the maps by multiplying each map by a gain and then adding the resulting map values. The gain for each map indicates the importance of that map (or behavior) in the current situation. The gain values are determined dynamically based on the state of the environment.

This work generalizes the idea of preferential behavior outputs from one dimensional action spaces to action spaces of arbitrary dimension. However, from the point of view of behavior design and command fusion and action selection this approach, like the approach in DAMN, is based on heuristics.

9.3 Action Voting [Hoff & Bekey 95]

An artificial neural network architecture is proposed in [18], where the emphasize is placed on making the action selection mechanism adaptive to short- and long-term environmental and goal changes. In each perception-action cycle each behavior votes for one action suitable from its point of view. Behavior priorities are modeled by adding to the action votes a positive bias value. Further, behaviors compete by inhibiting conflicting actions which essentially corresponds to voting against undesired actions. Inhibition is modeled as a nonlinear function of the strength of a behavior's action vote and the inhibition value for the desired action is set to 0 (i.e., no inhibition). The votes and inhibition values received from each behavior are summed for each action and the action with the highest value is then selected. In [18] each behavior is implemented by a neural network as explained in [14], however, this has no influence on their proposed action selection mechanism.

The main advantage of this approach lies in its ability to adaptively estimate each behavior's inhibition function based on feedback from each behavior's performance. In a nutshell a behavior's inhibition function is adjusted according to the following rules:

- When a behavior is not performing well and its current suggested action is not selected then its inhibition function is amplified to allow the behavior to compete and win.
- When the behavior's suggested action is selected and another behavior suffers, then the behavior's inhibition function is damped to give 'loosing' behaviors a chance to win.

The main results of this approach are that here it is demonstrated how the behavior inhibition function (corresponding to weights in DAMN) can be learned. For detail on how the inhibition function is updated see [18].

10 Fuzzy command fusion

In fuzzy behavior-based control, each behavior is synthesized by a rule-base controlled by an inference engine to produce a multivalued output. Figure 13, illustrates a simple example describing a collision avoidance behavior using only two rules. One rule recommends (expresses the desire) to turn away from a close obstacle and the other suggests to move straight forward if the obstacle is at a safe distance. Using standard fuzzy inferencing, e.g., max-prod in this example, the rules are combined into a multivalued output that encodes the (grade of) desirability of each action from the behavior's point of view.

Behaviors that compete for the control of the robot must be coordinated to resolve potential conflicts. Fuzzy behavior coordination is performed by combining the fuzzy outputs of the behaviors using an appropriate operator such as a triangular conorm (which corresponds to the fuzzy set union), e.g., the max operator. Then defuzzification (e.g., center of gravity, COG) is used to select a final crisp action ultimately used for control (see [56]). This scheme should lead to the selection of an action that somehow represents the consensus among the behaviors and thus comprises the action that best satisfies the decision objectives that they encode. However, in situations where there is conflict between the active behaviors, this approach leads to inappropriate results. One approach to dealing with conflicts is based on *context-dependent blending* [56], where context dependence is encoded using a number of fuzzy meta rules such as:

```
IF obstacle is close THEN avoid collisions
IF NOT(obstacle is close) THEN follow target
```

If the robot is close to an obstacle, then the contribution from the collision avoidance will be higher than the contribution from target following. This corresponds to weighted decision making, where the weight values are determined by the truth values of the rule antecedents and is used to scale the behaviors' fuzzy outputs accordingly. Another approach, used in the Saphira architecture [22], is to resolve conflicts by allowing higher priority behaviors to suppress/dominate behaviors with lower priorities. These approaches provide partial solutions to resolving conflict between behaviors in specific situations and systems.

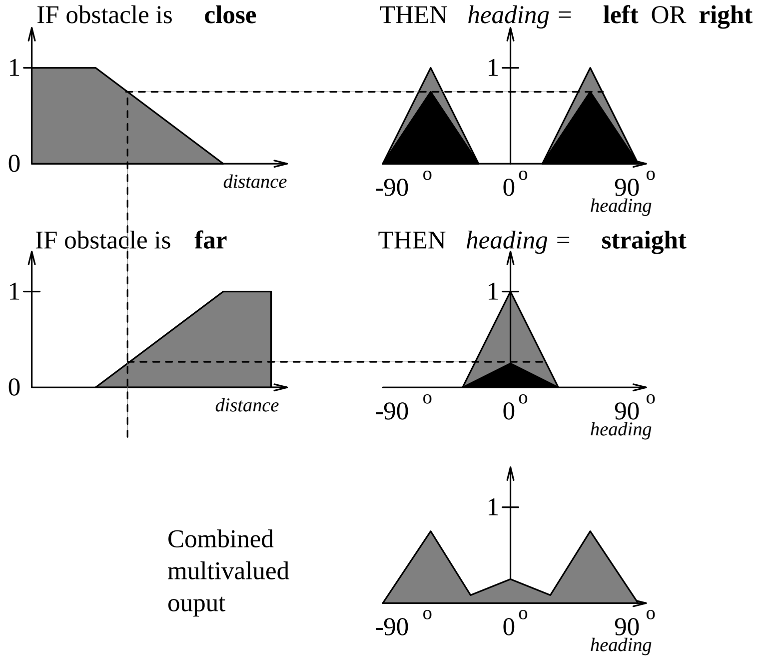

Figure 13: An example fuzzy rule-base to describe a collision avoidance fuzzy behavior.

10.1 Fuzzy/multivalued logic approach [Saffiotti, Konolige & Russpini 93]

The architecture described in this section is highly influenced by fuzzy logic concepts and many of its components are based on fuzzy formalism or multivalued logics [57]. This work describes not only the skeleton and specification of each component or layer of the architecture but it gives a remarkably detailed description of the formalism, mechanisms and design methodology of its components. The architecture is organized into 3 layers consisting of: control schemas, behavior schemas and planners, which interact with the agent as well as its environment at different levels of abstraction: signal / stimulus, intermediate (symbol-to-signal) and symbol respectively. Control schemas describe types of movements based on the agent's internal state and sensory input. Planners, at the other end, synthesize plans which are based on symbolic descriptions. Behavior schemas bridge the gap between the planners and control schemas by containing parts of both symbolic and signal levels of abstraction.

Control schemas are conceptualized as mappings from a set S of states to preferences on a set A of actions to perform. Formally the mapping

$$D : S \times A \rightarrow [0, 1], \tag{1}$$

measures the desirability, $D(s, a)$, of performing control action a in situation s. In practice each control schema is implemented by a set of fuzzy IF-THEN rules. Using standard fuzzy reasoning, the outputs of the rules are combined into a membership function reflecting the desirability D of the implemented control schema. The resulting membership function is defuzzyfied to compute a single control value. Further, control schemas can be composed using operators from fuzzy logic, AND, OR etc. For example, using the AND operator two control schemas can be blended into a control schema that, under certain conditions, will satisfy the desirability functions of both control schemas. While defuzzification can produce appropriate results in some cases it may, however, in cases where there is a conflict between behaviors, defuzzification might produce inappropriate results (see section 10.2 on a discussion of pros and cons of several defuzzification methods). E.g., if a behavior desires to move left and another desires to move right then defuzzification will result in moving straight ahead which does not confirm to the desires of any of the behaviors. In order to alleviate this problem Saffiotti associates each control schema to a context, which defines the conditions under which the control scheme is applicable. Contextual conditions are implemented by a set of (meta) rules that activate control schemas

$$\text{IF context} = A \text{ THEN activate\_control\_scheme } C,$$

where A is a fuzzy term representing the actual context. This fuzzy rule activates control scheme C to a level determined by the degree of truth of the rule's antecedent. The activation level is used to weight the desirability function of the control scheme. In this way control schemas can be designed, which can furnish certain actions such as walk, jump and sit. Control schemas operate on internal variables such as movement direction and speed.

Behavior schemas are the structures that tie certain action patterns, implemented by control schemas, with certain stimuli in an environment. In this way stimulus-response behaviors can be formulated that accomplish specific goals such as obstacle avoidance. A behavior is described by a triple

$$B = (C, D, O), \tag{2}$$

where D is a control schema for a specific type of movement, O is a set of object descriptors with respect to which the movement should be performed and C is a contextual condition, which describes the context of applicability of the movement. Object descriptors are models of objects in the real world and they cause a behavior to act with respect to the external world as opposed to a control schema that acts upon internal variables. The object descriptors function similar to perceptual markers introduced in [2], where they are updated each time the perceptual system observes the corresponding objects in the real world in which way a closed-loop is formed. Deliberative goal achievement in this architecture is handled by inserting certain artifacts in the perceptual space associated with the behaviors that are designed

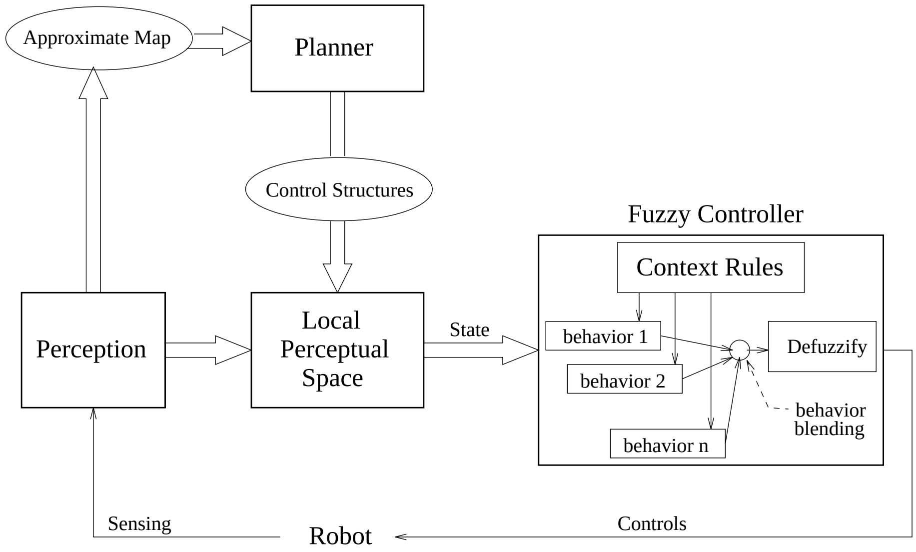

Figure 14: Flakey's Architecture, adapted from [58].

to respond appropriately to these artifacts. These artifacts are denoted control structures and they represent goals in this framework. For example, a 'control-point' is an artifact representing a position to be reached and Go-To-CP is a fuzzy behavior implementing how the robot should react to reach this position.

Once a set of 'primitive' behaviors have been formulated in this way, they can be composed into more complex behaviors using the composition technique of control schemas described earlier in this section. Saffiotti defines three composition operators for conjunction, blending and chaining of behaviors, not explained here (see [57] for details). Further behaviors have specifications in terms of preconditions (context of applicability) and postconditions (goal accomplishment), which can be used by traditional planners to construct plans for achieving goals. Goals, in this framework, are defined as predicates on all the possible executions of a behavior. The formulation of behaviors and goals can be used to give formal proofs of various characteristics of combined behaviors and determine what their combined effect will be in terms of goal accomplishment. Planners, from traditional AI, can then use the specifications of behaviors in terms of preconditions and postconditions and the machinery for reasoning about composed behaviors to construct plans that accomplish given goals.

One major disadvantage of fuzzy approaches is that there are no clear guidelines to fine tune the fuzzy membership functions, which renders the fine tuning process tedious. However, learning techniques are being developed that can help in this process. This architecture has successfully been deployed for control of an in-door mobile robot system, Flakey. Flakey finished second in the AAAI robot competition of 1993 [41]. A similar approach is presented in [66], where behaviors are synthesized by a set of fuzzy IF-THEN rules and the activation level of each behavior is determined by another layer of rules.

10.2 Fuzzy DAMN [Yen & Pfluger 95]

Inspired by the work of Payton and Rosenblatt [44] but not satisfied with the implicit heuristic treatment of their ideas, [70] propose a fuzzy logic formulation of the command fusion approach. The ideas presented by Yen and Pfluger were already put forward in [57] without a direct link to Payton and Rosenblatt's approach. Yen and Pfluger make this link explicit.

Yen and Pfluger realize that certain fundamental issues in the DAMN architecture bear a great resemblance to concepts in fuzzy logic. Thus they establish the relation between the concepts introduced by Payton and Rosenblatt and fuzzy set theory and control. Specifically, first the outputs of behaviors can be cast as discrete membership functions over the set of possible actions. Second, the command fusion based on weighted summation in DAMN can be replaced with standard fuzzy inferencing methods (e.g., max-min, weighted-sum) in its, here, fuzzy formulation. Third, maximum-vote action selection can be replaced by defuzzification techniques in fuzzy logic (e.g., center of gravity defuzzification). Further, Yen and Pfluger propose to implement each behavior using a set of fuzzy rules (see figure 15).

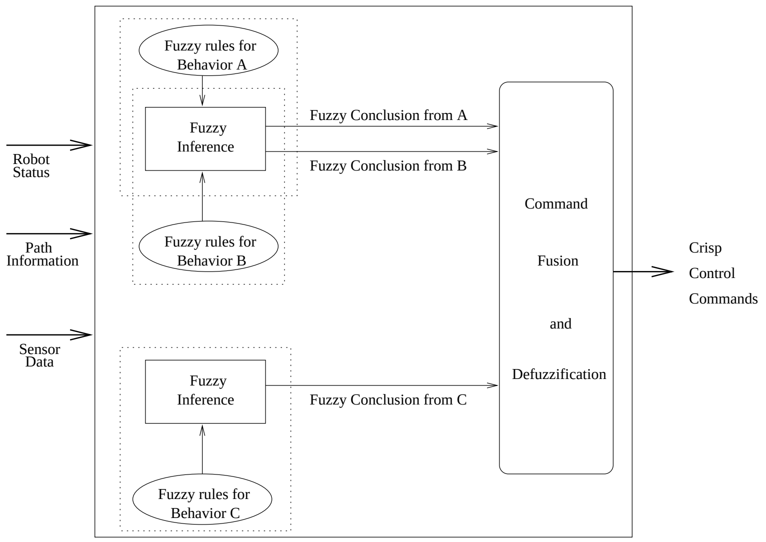

Figure 15: The DAMN architecture based on fuzzy logic, adapted from [70].

The fuzzy logic formulation of Payton and Rosenblatt's approach introduces useful guidelines in design of behavior-based control systems. However, command fusion and action selection using standard fuzzy inferencing and defuzzification techniques respectively can have certain disadvantages as demonstrated in [70]. Using the example in figure 16, it is illustrated that the center of gravity, COG, (also known as the center of area) technique produces inappropriate results.

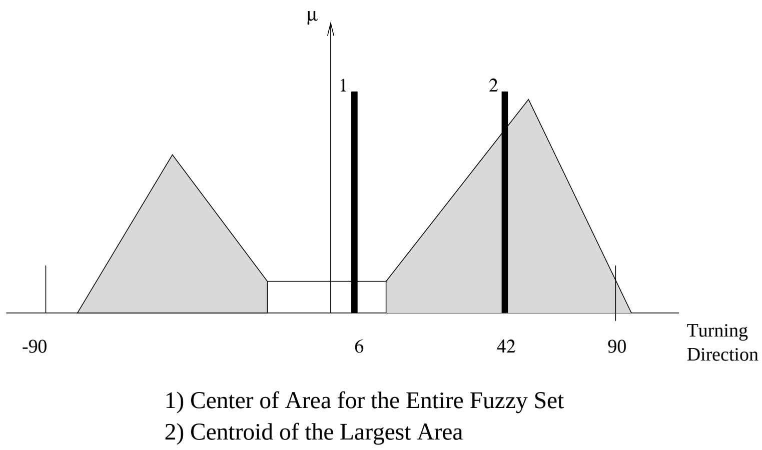

Figure 16: Examples of defuzzification techniques [70].

The COG defuzzification technique in the example depicted in figure 16 selects an action that does not fulfill any behavior (here there are two behaviors, obstacle avoidance and path following). The authors thus propose another defuzzification technique called Centroid of Largest Area, CLA. This technique selects the command that corresponds to the largest area of a multi-modal fuzzy set. For this particular example this technique produces an appropriate result, see figure 16, however, this does not mean that CLA will always provide 'good' actions. For instance, in the example shown in figure 17 CLA will produce an action that is not acceptable.

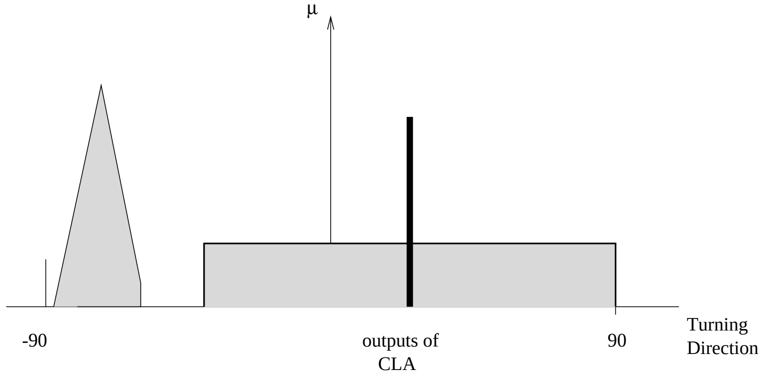

Figure 17: Examples of the CLA defuzzification technique, where the produced output is not appropriate.

11 Multiple objective behavior coordination

The pholosophy behind this approach is that, in real-world robotics it is not feasible to aim for *optimal* control even if an optimal strategy should exist, because unstructured environments are unpredictable, robotic sensors and actuators are imperfect, and computational and other resources (such as response time) are sparse. It is thus more appropriate to search for control strategies that in practice are *efficient* and *sufficiently satisfying*, i.e., 'good enough' rather than optimal. This in fact has been an implicit agenda in behavior-based robotics all along. However, the view taken by this approach is that only by having this as an explicit goal can we hope to produce efficient (not just good enough) behavior. In the behavior-based approach there are no centralized planning modules that search for plans with some optimality criteria (e.g., shortest path, safest path, or minimal power consumption). Rather, control in behavior-based robotics is shared between a set of sensory-motor units, called behaviors, where each behavior is responsible for one specific aspect of control (such as obstacle avoidance, target following, door traversal etc.) in isolation from the other behaviors. Note that a solution that is simultaneously optimal with respect to all the behaviors might not exist due to conflict between behavioral objectives. The optimization of one behavior's objective might be associated with a simultaneous deterioration of other behaviors' objectives. Thus a major issue in behavior-based control is to formulate strategies that can guarantee a suitable trade-off between a multitude of (potentially) conflicting behavioral objectives.

The view promoted by this approach is that behavior-based control strategies should *explicitly* define a notion of optimality that is both suitable and practically feasible for the kind of problems that it is trying to solve. Only then will it be possible to specify and design systems that can generate efficient and 'good enough' behavior. This approach proposes a notion of optimality for behavior-based control which is then used to formulate a class of action selection mechanisms called Multiple Objective Behavior Coordination (or Action Selection) Mechanisms.

Multiple objective action selection mechanisms incorporate an explicit formulation of control objectives into the system description. By so doing it is possible to evaluate each alternative action against all objective functions simultaneously so to find the action that satisfies the objectives as good as possible. Acknowledging the fact that optimal actions might not exist due to conflicting objectives, MOASMs seek for *satisficing* actions, i.e., actions that *satisfy* each objective function to a *sufficient* degree. I postulate that satisficing actions compromise 'good enough' actions in the action selection problem.

It has been demonstrated that multiple objective decision theory provides a suitable formalism to encompass these ideas into a framework for action selection in behavior-based control. In a nutshell, each behavior calculates an objective function over a set of permissible actions. The action that maximizes the objective function corresponds to the action which best satisfies that objective. Multiple behaviors are blended into a single more complex behavior that seeks to select the action that satisfies simultaneously all the objectives as well as possible. In this approach action selection is comprised of generating and then selecting a (feasible) set of satisficing solutions among the set of efficient solutions known as Pareto-optimal solutions. The theory of multiple objective decision making (MODM) proposes methods for efficient generation of such solutions [12, 63].

11.1 Multiple objective behavior coordination

Multiple objective behavior coordination is based on multiple objective decision theory [12]. Since multiple objective decision theory has many concepts that have synonymous counterparts in fuzzy logic, we only introduce aspects of the theory that support the discussion of MOBC.

11.2 Objectives and behaviors

Objectives define the standards against which the quality of a given alternative may be evaluated. An objective function, o, is thus a mapping that assigns to each alternative, $x \in X$, a value reflecting that alternative's desirability. Multiple objective decision making (MODM) provides approaches for making decisions in complex situations where more than one decision objective should be considered.

In multiple objective behavior coordination, a behavior, b, is formalized as an objective function that defines a mapping from an action space, X, to the interval $[0, 1]$: $b : X \to [0, 1]$. The action space, $X = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$, is defined to be a finite set of alternative actions. Behaviors in this framework are equivalent to fuzzy behaviors because objective functions encode the semantics of fuzzy membership functions.

11.3 Multi-objective decision problems and behavior coordination

Once the link between behaviors and objective functions is established, we can formulate the coordination of multiple behaviors as a multiple objective decision problem that seeks to *optimize a number of objective functions simultaneously*, where the objective functions correspond to behavior multivalued outputs. From the set of alternative actions X, an action $\hat{\mathbf{x}} \in X$ should be chosen that best satisfies all objective functions. This is formulated as a multiple objective decision problem in the following way:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{argmax} \qquad [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})] \qquad (3)$$
$$\text{subject to} \qquad \mathbf{x} \in X,$$

where $o_i(\mathbf{x})$ for $i = 1, ..., n$ are the objective functions estimated by corresponding behaviors b_i and $\mathbf{x}$ is an N-dimensional control vector $(x_1, x_2, ..., x_N)$. Further $X \subseteq R^N$ is the constrained action space[3]. The solution $\hat{\mathbf{x}}$ to equation (3) corresponds to the selected action that best satisfies the corresponding behaviors.

11.4 Solving multiple objective decision problems

There are a number of methods described in the literature for solving multi-objective decision problems each with certain advantages and disadvantages [46, 12]. Rather than merely giving a description of these methods, we highlight the main characteristics of multiple objective decision making that make it suitable for behavior coordination. Thus, we start by describing the types of solutions to a MODM problem.

A solution is *optimal* if it results in the maximum value of each of the objective functions simultaneously. That is, $\hat{\mathbf{x}}$ is an optimal solution to equation (3) iff $\mathbf{o}(\hat{\mathbf{x}}) \geq \mathbf{o}(\mathbf{x})$ for all $\mathbf{x} \in X$, where $\mathbf{o}(\mathbf{x}) = [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})]$.

An important characteristic of MODM problems is that the multiple objectives, $o_1, ..., o_n$, can be competitive and conflicting, i.e., the improvement in one can be associated with a deterioration in another. Hence optimal solutions might not exist and thus the concept of optimality should be modified to endow similar and useful concepts in MODM problems. In practice we can only hope for solutions that *satisfy* each of the objectives to a *sufficient* degree. This idea is encompassed in the following definition.

Definition 11.1 (A Satisficing Solution) *A satisficing solution of Simon[62] extended to a multiple objective decision problem is a reduced subset of the feasible set, which exceeds all of the aspiration levels (goals) of each objective.*

Simon points out that it might be impractical to search for an optimal solution and argues that in practice it is more efficient to look for solutions that are 'good enough'. A satisficing solution corresponds to the behavior of a decision maker who looks for 'good enough' solutions. Another useful notion of optimality in the context of multiple objective decision making is *Pareto-optimality.*

Definition 11.2 (A Pareto-Optimal Solution) $\mathbf{x}^*$ *is a Pareto-optimal solution to a multiple objective decision problem iff there does not exist any* $\mathbf{x} \in X$ *such that* $o_i(\mathbf{x}^*) \leq o_i(\mathbf{x})$ *for all* i *and* $o_j(\mathbf{x}^*) < o_j(\mathbf{x})$ *for at least one* j.

Intuitively, a Pareto-optimal solution is one in which no other solution can improve one objective function without a simultaneous deterioration of at least one of

[3]The action space is constrained due to various physical and geometrical properties, such as a set of non-holonomic constraints imposed by the mechanical construction of the robot, into a subset of *feasible* (e.g., physically possible) actions.

the other objectives. A Pareto-optimal solution is also known as a nondominated solution, a non-inferior solution, or an efficient solution.

11.5 Multi-objective behavior coordination

It is not so much the methods for generating solutions but the notions of Pareto-optimal and satisficing that play a significant role in formulating behavior coordination mechanisms with useful properties. See [46] for techniques for generating PO and satisficing solutions for multi-objective decision problems. Based on the definitions of Pareto-optimality and satisficing, we propose a behavior coordination scheme that proceeds in the following step-wise fashion:

1. **Feasible actions**. Based on physical and other hard constraints find the set of feasible and permissable alternatives
2. **Pareto-optimal actions**. Based on objective criteria find the set of Pareto-optimal actions and reject others to reduce significantly the set of alternatives
3. **Satisficing actions**. Incorporate subjective knowledge (weights, priorities, goals, aspirations levels) to find the set of satisficing alternatives
4. **Most preferred action**. Use other criteria to finally select the most preferred action.

This approach systematically narrows down the set of possible solutions, by including increasingly more decision criteria. First, the set of hard constraints given by the nature of the task and physical limitations of the robot (e.g., limited acceleration) are used to exclude the set of non-feasible actions. Then, based only on the objective criteria, we can exclude the set of inferior or dominated solutions simply because they are not efficient when compared to the Pareto-optimal solutions. By incorporating some subjective preference structure in terms of weights, priorities etc. we can then select, out of the Pareto-optimal solutions, the set of satisficing, i.e., sufficiently satisfying, solutions. Finally, other criteria or further subjective knowledge can be included to make a final selection. In this way alternatives are evaluated against the behavioral objectives and ranked in terms of permissable, Pareto-optimal, satisficing and preferred actions.

11.6 Decision-theoretic fuzzy behavior coordination

Rather than letting the behavior coordination depend on the operators provided by the underlying fuzzy mechanism, multiple objective behavior coordination is first and foremost concerned with the quality of the selected actions (in terms of Pareto-optimality and Satisficing). However, as it should be evident by now, the (input/output) mechanics of multiple objective behavior coordination are not different from the fuzzy approaches, hence they can be interchanged.

Further, several of the approaches for generating solutions (Pareto-optimal and satisficing) to multiple objective decision problems (see [46]) use weight values . The weight values have the same role as in context-dependent blending, and determine the degree of contribution for each behavior. Thus, also with respect to behavior blending the fuzzy and multiple objective approaches are compatible and thus can be combined. One straight-forward way to combine both approaches is to:

1. Specify each behavior using a fuzzy rule-base, and generate multivalued behavioral outputs using fuzzy inferencing,
2. Use fuzzy context rules to generate weight values, and
3. Coordinate the behaviors by generating Pareto-optimal and satisficing solutions using multiple objective decision making techniques.

In this way we combine the advantages of approximate reasoning and ease of behavior specification of fuzzy logic with the qualitative and formal behavior coordination capabilities of multiple objective decision making. This approach was introduced in [49] and has been shown in [48] to have certain advantages over standard fuzzy inferencing techniques. In [50] the approach has successfully been applied to control of single as well as multiple robots.

12 Superposition based command fusion

12.1 Potential fields [Khatib 86]

The potential-field approach introduced by [21] is an approach to motion planning where the robot, which is represented as a point in configuration space, moves under the influence of an artificial potential field produced by an attractive force at the goal configuration and repulsive forces at the obstacles. Action selection in this terminology corresponds to at each configuration to move in the direction indicated by the negated gradient of the total potential denoted U.

Thus for obstacle avoidance, the potential function U is constructed as the sum of two more elementary potential functions:

$$U(q) = U_{att}(q) + U_{rep}(q) \tag{4}$$

where U_{att} is the attractive potential and U_{rep} is a repulsive potential associated with obstacles. The force $\vec{F} = -\vec{\nabla} U$ corresponds to the most promising direction to move in. From equation (4) it is seen that $\vec{F}$ is the sum of two forces: attractive $\vec{F}_{att} = -\vec{\nabla} U_{att}$ and repulsive forces $\vec{F}_{rep} = -\vec{\nabla} U_{rep}$.

There are many ways to define the elementary potential functions. For example [29] defines the attractive potential field U_{att} as a parabolic well:

$$U_{att}(q) = \frac{1}{2}\xi ||q - q_{goal}||^2, \tag{5}$$

where $\xi > 0$ is a scaling factor and q_{goal} denotes the goal configuration. Figure 18(b) is a plot of U_{att} for a two dimensional case, i.e., translation in x and y, depicted in figure 18(a).

The basic idea behind the definition of repulsive potentials is to create a force that will push the robot away from obstacles. [29] defines a repulsive potential function as follows:

$$U_{rep}(q) = \begin{cases} \frac{1}{2}\eta(\frac{1}{\rho(q)} - \frac{1}{\rho_0})^2 & \text{if } \rho(q) \leq \rho_o, \\ 0 & \text{if } \rho(q) > \rho_0, \end{cases} \tag{6}$$

where $\eta > 0$ is a scaling factor, $\rho(q)$ denotes the distance from q to the closest obstacle and $\rho_0 > 0$ is the distance within which the robot is under the influence of the repulsive forces. Figure 18(c) is a plot of the repulsive potential field function for the example. Figure 18(d) is a plot of the total potential field function U and its negated gradient is depicted in figure 18(e). Note that regardless of the start configuration, q_{start}, following the direction of the force $\vec{F} = -\vec{\nabla}U_{att} - \vec{\nabla}U_{rep}$ will bring the robot to the goal configuration q_{goal}.

The potential fields approach has been successfully used for planning collision-free path for robotic arms [21] as well as for mobile robot platforms [29, 13]. However, the usual formulations of potential-field methods do not preclude the occurrences of minima other than the goal. Thus, in the steepest-decent search the robot can fall into these minima and achieve a stable configuration short of the goal state. In [13] the application of *harmonic functions* to potential-field path-planning is described. One of the properties of this approach is that it guarantees the exclusion of spurious local minima within the solution region. In practice, however, due to quantization limitations and numerical instabilities, local minima may occur. Further, the calculation of the harmonic function can be rather time-consuming for relatively large configuration spaces.

In its original formulation, the potential-field method had no explicit links to action selection problem for behavior-based control. This relation was established in the motor schemas approach to behavior coordination and action selection [4].

12.2 Motor schemas [Arkin 87]

This approach is based on the concepts of *schemas* and *potential fields*. Schemas originate in psychology and neurology and schema theory, introduced by Michael Arbib, describes models which explain how the mental process organizes perception and coordinates responses to environmental stimuli (see [4] for a discussion on schema

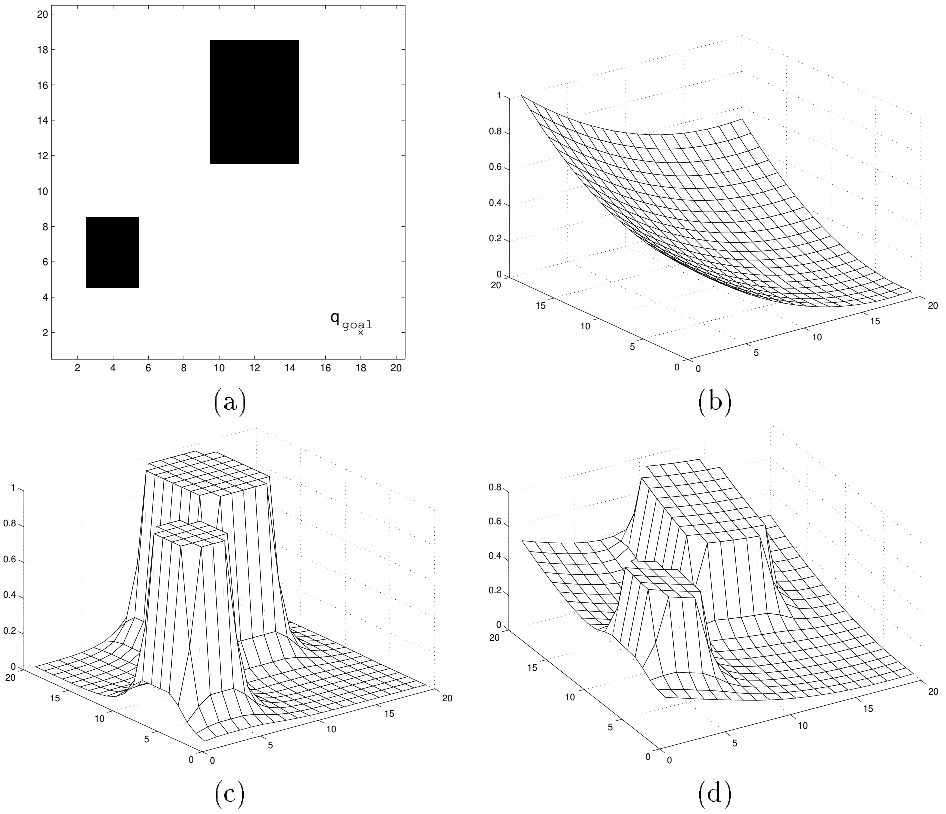

Figure 18: (a) The configuration of obstacles (black) in the map. (b) The attractive potential associated with the goal configuration. (c) The repulsive potential associated with obstacles. (d) The sum of the attractive and repulsive potentials.

theory and its links to robotics). Schemas are parameterized potential functions[4], which give a generic specification of independent process, that are specialized for a specific task and domain. There are two types of schemas: *motor schemas*, which are concerned with control of actuators; and *perceptual schemas*, which are concerned with goal-directed sensing of features in the environment. From a software engineering point of view, schemas are well defined entities that are "easy to comprehend and handle" [4]. As Arkin puts it "*The task of robot programming is fundamentally simplified through the use of a divide and conquer strategy*". In particular, perceptual schemas allow an agent to direct its sensing resources to processing of information that is pertinent for a given task. For example given an obstacle position, an obstacle avoidance motor schema can control the robot away from the obstacle. In order to detect an obstacle or obstacles, one or several obstacle detection perceptual

[4]At run time the schemas are instantiated for a particular task.

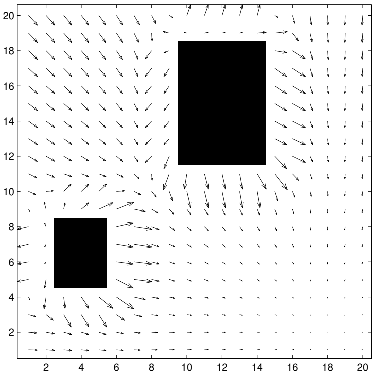

(e)

Figure 18: (continued) (e) The negated gradient vectors.

schemas can be instantiated to keep track of obstacles and feed the obstacle positions to associated obstacle avoidance motor schemas.

The system architecture is depicted in figure 19. During operation, the invoked schema instances (SI) are composed and coordinated to accomplish specified system goals. Composition is used to combine several independent motor schemas into a more complex schema that can achieve all their goals in a consistent manner. Each motor schema instance is designed to generate a vector which encodes the direction and intensity (e.g., velocity) of motor action to perform for achieving the motor schema's goals. Thus each motor schema calculates a potential-field, not for the entire configuration space, but only for the current configuration of the robot. The vectors generated by each SI are added to generate a combined motor action. Further since motor schemas can have different importance, each SI has a gain value that is multiplied by the its generated vector. The gain values are usually determined empirically for a given task and environment. Case-based learning mechanisms have been studied, which adapt the gain values (and other system parameters) according to the robot context [51].

Temporal sequencing (explained in section 7.2) is another coordination mechanism used in the architecture depicted in figure 19.

The schema based approach has been used for construction of mobile robot systems as early as 1987. Autonomous robot architecture (AuRA) was a three layered architecture consisting of a mission planner, navigator and pilot. The mission planner performed long term planning, the navigator produced a piecewise path and the pilot coordinated a set of schemas to furnish the actions that implemented the goals set by the navigator. In the present time, researchers at the Georgia Institute of Technology have made considerable advance in developing schema based mobile robotic systems. In particular the team of mobile robots, Io, Ganymede and Callisto

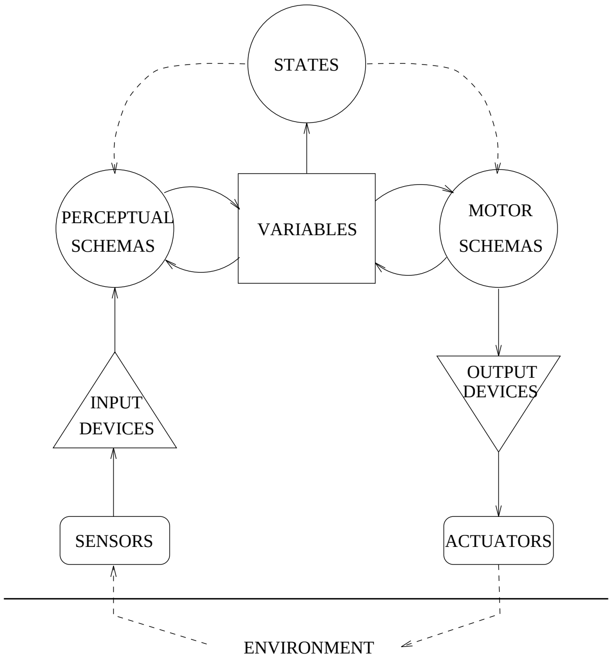

Figure 19: The Control Architecture of the Motor Schemas Approach. From [7]

have been successfully deployed in cooperative activities. They won the AAAI 1994 robot competition by collecting more cans in 10 minutes than any other contestant.

12.3 Dynamical systems approach [Schöner 92]

A fundamentally different approach to action selection was proposed in [59] based on dynamical systems theory. This approach proceeds in the following way:

1. **Behavior generation**: In this approach each behavior is described by a differential equation expressing the relation between some *behavioral variable* and its time derivative given by the following general form: $\dot{\phi} = f(\phi)$. ϕ corresponds to the behavioral variables such as velocity or robot heading. Further, *task constraints* from the environment are used to describe the function f. These constraints are modeled as two types of primitive forces: *attractors* and *repellors*; that are used to describe more complex behaviors.
2. **Behavior composition**: The contributions from the behaviors are combined additively by means of superposition:

$$\dot{\phi} = \sum_i f_i(\phi) + noise. \tag{7}$$

3. **Action selection**: Behavior is generated by evolving in time the solutions of the dynamical system that governs the behavioral variables.

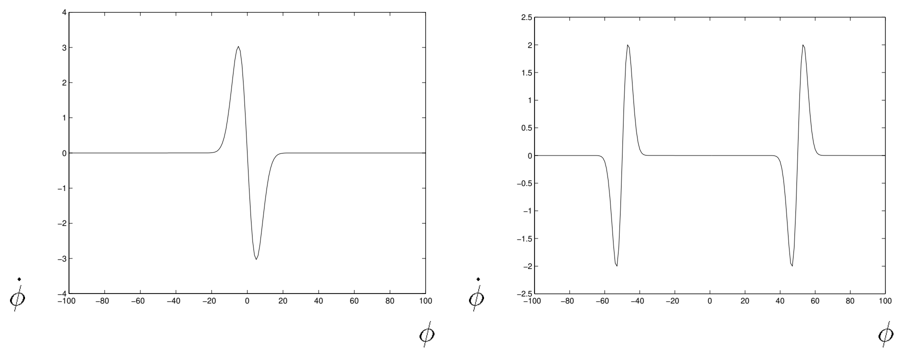

Figure 20: (left) Phaseplot of a simple dynamic system with one attractor and (right) two repellors.

For instance a behavior for moving to a target with heading ϕ_{tar} can be described by:

$$\dot{\phi}_{tar} = \lambda(\phi_{tar} - \phi)e^{-\frac{(\phi_{tar}-\phi)^2}{2\sigma^2}}, \ \lambda > 0, \tag{8}$$

where the control variable ϕ is the heading direction of the robot. λ is the strength of the attractor and can be made a function of the distance of the target to the robot. σ defines the width of the attractor window and can be a function of the robots diameter [64]. Equation (8) for $\lambda < 0$ defines a repellor and can be used to describe an obstacle avoidance behavior.

$$\dot{\phi}_{obs} = \sum_i \lambda_i(\phi_i - \phi)e^{-\frac{(\phi_i-\phi)^2}{2\sigma^2}}, \ \lambda_i < 0, \tag{9}$$

where ϕ_i are the directions to the obstacles. The plots in figure 20 depict simple dynamic systems, one with a single attractor and the other with two repellors. The sum of these systems is plotted in figure 21.

The behaviors of target following described in equation (8) and obstacle avoidance in equation (9) are combined in the following way (see figure 21):

$$\dot{\phi} = \dot{\phi}_{tar} + \dot{\phi}_{obs} + noise. \tag{10}$$

Desired behaviors of the system are then given by the set of stable solutions or fixed points of equation (10). The reason for adding a noise component to the behavioral dynamics is to ensure that the system does not fall into a stable attractor, since this will prevent any evolution over time. A major advantage of this formal approach is not only the formal basis for design and description of behaviors but also that it allows for a theoretical analysis of system characteristics such as stability.

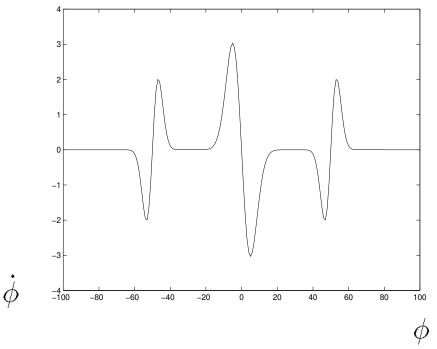

Figure 21: Phaseplot of a simple dynamic system with one attractor and two repellors.

One problem that can occur when repellors and attractors fall close to each other is the generation of a spurious attractor that can bring the robot in undesired situations. To handle such cases an extension to the dynamical approach is proposed by [28]. Basically it is proposed to modify the contribution from each behavior depending on the situation at hand. A weighting factor is associated with each behavior that determines its contribution to the behavior dynamics:

$$\dot{\phi} = |w_{tar}|\dot{\phi}_{tar} + |w_{obs}|\dot{\phi}_{obs} + noise. \tag{11}$$

Based on the current context appropriate weights are calculated for each behavior. Formally, these weights are calculated as solutions to a *competitive dynamics* given by the following equation [28].

$$\dot{w} = \alpha_i w_i (1 - w_i^2) - \sum_{j \neq i} \gamma_{j,i} w_j^2 w_i + noise. \tag{12}$$

The 'current context' is described in terms of the *competitive advantage* and *competitive interaction* given by the parameters, α_i and $\gamma_{j,i}$ respectively. Informally, competitive advantage defines the priority of behavior i and competitive interaction is a measure of conflict between behaviors j and i. Geometrically, conflict can be expressed as the distance, in the dimension of a behavioral variable, between an attractor and a repellor. E.g., maximum conflict is introduced if a repellor and an attractor fall on each other. Thus by an analysis of the system the parameters α_i

and $\gamma_{j,i}$ are designed to produced a desired behavior: what behavior is preferred, given by α_i, when a conflict arises, given by $\gamma_{j,i}$.

Using this technique [28] have demonstrated that behavior switching and averaging can be achieved. Behavior switching corresponds to a stable solution to the competitive dynamics where the weights for one of the behaviors is 1 and for others 0. Another stable solution produces an averaging behavior where all behaviors are active.

13 Discussion

In the previous section we have listed and in some detail described a set of approaches to the action selection problem. We realize that this list is incomplete, however, the work that we have selected to described, we believe, includes relevant representatives of the activities in this area. Various important issues for action selection mechanisms are discussed here.

13.1 Single-valued vs. multi-valued behaviors

Traditionally behaviors typically produce a single-valued output that from the behavior's point of view is the optimal action to perform. A problem that is commonly faced in multi-behavior systems is that the generated actions by different behaviors may be contradictory. Thus a scheme should be used to resolve conflict situations. If the behaviors are in conflict then it is difficult to make good decisions based on single-valued behaviors. E.g., in the motor schemas [4] each behavior produces a single-valued vector output, the selected action is then the weighted average of these vectors. The produced control vector of this strategy is, however, not necessarily appropriate because it might produce a result that does not satisfy any of the behaviors.

From the point of view of a single behavior its produced output is the best output. In composing behaviors, i.e., deciding what action to take based on the outputs of each behavior, we have to decide what action is the most appropriate. This could be formulated as finding the action that makes the best tradeoff between the desired actions of each behavior. This, however, is an ill-posed problem if the decision should be based only on the single-valued outputs generated by each behavior.

Rosenblatt and Payton [53] demonstrated that, if each behavior generates a multi-valued output a more moderate composition scheme could be constructed. Multi-valued outputs seem to make an appropriate basis for communication among the behaviors. The outputs generated by the fuzzy behaviors of the fuzzy approaches can also be considered multivalued, since their outputs are fuzzy membership functions. It seems that multivalued outputs define an appropriate interface and means of communication between the modules. In [43] Payton explains why multivalued outputs portray an appropriate interface between system components. One of his points is that the amount of loss of information is reduced thus it defines a better basis for communication. One problem, however, encountered by Rosenblatt and Payton is that the semantics of the behavior outputs or the action selection process is not very clear.

13.2 Behavior and action selection semantics

The basic ideas of the approach presented by Rosenblatt and Payton were further refined and formulated in terms of fuzzy set theory and control. The fuzzy command fusion methods describe a behavior (output) as a fuzzy member-ship function and reflects the (grade of) desirability of each action from that behavior's point of view. Then the selected action is interpreted as the action that satisfies as good as possible the behaviors. However, this point is not clear and as pointed out in [70] standard defuzzification methods can lead to actions that are not good for any behavior at all. Thus other defuzzification methods are suggested. In particular, multiple objective behavior coordination proposes approaches to action selection that can replace standard fuzzy inferencing with multiple objective decision making while maintaining the advantages of fuzzy behavior coordination. This combined multi-objective/fuzzy approach is known as decision-theoretic fuzzy behavior coordination.

A formal approach to action selection enables a theoretical analysis of action selection and characterization of the selected actions.

13.3 Formalism & performance guarantees

One of the main challenges of mobile robotics is to devise a theoretical framework for analysis, design and synthesis of systems that have specified properties, performance and behavior. Formal approaches add the power of prediction to the framework. This is an important characteristic, since we would be able to design systems using specification and prove various features about the system. Then the process of design and construction of mobile robot systems would be much similar to the discipline of, e.g., classical control theory. However, we realize that in order to allow the formalism to capture the essence of a system one should make certain approximations and simplifications.

In the dynamical systems approach, the fuzzy approach [57] and the DES formalism [23] it is seen that stability, formal proofs of task achievement and controllability guarantees, respectively, can be given when the approach is based on a formal framework.

Multiple objective behavior coordination mechanisms provide formal tools for characterizing action selection in terms of Pareto-optimal and satisficing solutions. Satisficing behavior is proposed as an alternative to optimal action selection because in practice optimal solutions might not exist due to potential conflicts between the objectives of competitive behaviors. By having an explicit formulation of objective functions, this method allows for a qualitative evaluation of each alternative action against the decision objectives.

13.4 Behavior synthesis

In the philosophy of behavior-based system synthesis, a complex control task is divided into a set of simpler control problems which then are implemented by behaviors. In the catalogue of approaches presented in this paper a number of different behavior synthesis methods have been proposed. Basically these approaches can be divided into *preprogrammed* and *learned* behaviors. Hand-coding of behaviors can be cumbersome, why it is desired to 'automatically' generate or learn behaviors. The main advantage of learned behaviors is demonstrated by, e.g., [35, 36]. However, a problem seems to be that in many learning approaches to behavior synthesis it is as hard to design good reinforcement strategies as it is to hand-code behaviors.

A different approach to implementation of behaviors is seen in the fuzzy approaches, where behaviors are represented by a set of if-then rules together with a inference engine, in this case a fuzzy reasoning mechanism. In the motor schemas approach, behaviors are synthesized using primitive types representing attractive and repulsive forces. Attractive forces correspond to desired behavior while repulsive forces correspond to behavior that should be avoided. A similar philosophy is used in the action maps approach used in SAMBA and the dynamical systems approach. There is, however, a fundamental difference between the dynamical systems approach and the motor schemas. And that is, in the dynamical systems approach the system is always in an attractor whereas in the motor schemas approach an attractor represents a terminal configuration.

Finally, in arbitration methods high-level behaviors are composed by coordinating the activation of lower-level behaviors. In many of the approaches in this category, finite-state machines are considered a useful tool for generating complex behavior.

13.5 Centralized vs. distributed control

The presented approaches have different views on whether control should be centralized or distributed. Most of the approaches have centralized control, however, Maes' approach is an example of a system, where control is distributed among all behaviors and there is no need for a bureaucratic module. Distributed approaches tend to be more robust to failures of single system components and enable graceful degradation. Further from a software engineering point of view, behavior-based systems enable modular system development. These are good reasons to choose a decentralized control scheme, but distributed systems add the overhead of coordination, synchronization and communication between system components, which can be overcome by powerful software engineering tools.

13.6 Planning vs. reactivity

The description of action selection mechanisms presented in this paper is focused on behavior-based robot control. It is important to note that for effective use of sparse system resources and efficient task achievement a robot must plan ahead of time. This is a view shared by many of the presented approaches, but how to integrate planning into such reactive systems is not obvious.

For the most part planning and reactivity are integrated into hybrid control to complement each others shortcomings. The basic idea is to use planners for what they are good at and are able to handle and use reactive systems for what they are good at.

As stated in the Introduction, traditional planners assume complete knowledge of the environment and the agent's sensing and actuating capabilities. Experience shows that planners can effectively be used to synthesis plans based on knowledge which is known to be complete (for the purpose). In [68] Turner et al. put it in the following way:

> *"Symbolic AI works best for well-defined, representable problems. On the other hand, what practical autonomy needs are systems which exhibit robust behaviors in ill-defined, poorly-constrained, and often unfriendly environments."*

Knowledge which remains the same, throughout the lifetime of the agent, and which does not change considerably is an appropriate candidate for use in traditional planners. Most environments contain structures (and relations between these) that remain static during (and after) the lifetime of an agent and which can be included in a static knowledge base of traditional planners. The location, size and relationship between rooms, walls and doors are examples of knowledge that do not change considerably. In [47] such global knowledge is used successfully to plan a sequence of actions to bring a robot from one room to another in an indoor office environment. It is rather easy to encode and incorporate such knowledge into the system and it simplifies the task significantly. Use of such knowledge enables the robot to use strategies appropriate for the context it is in. Such contextual information can be encoded in the knowledge base of the traditional planner, thus, the system does not have to detect what context it is in (as in case-based learning techniques [67]).

On the other hand, traditional planning approaches are not good at handling unpredictable situations, where we rather use reactive system components. By sensing repeatedly and acting accordingly, reactive system components enable the agent to detect and handle changes in its immediate surroundings. In this way the agent is given the ability to plan ahead of time in face of incomplete knowledge of the environment.

Further traditional planners usually have unbounded time constraints, which might slow down the system if not coupled appropriately with the reactive parts.

13.7 Integration of planning and reactivity

A review of work done in hybrid system architectures reveals that the interfacing between traditional planners and reactive modules can be achieved in various interesting ways. However, one thing that there seems to be common consensus about is that *planners should guide,* **not control** *reactive behaviors.* This corresponds to utilizing the *plan-as-resource* view, as opposed to the *plan-as-program* view. These fundamentally different views of plans have been promoted in [1].

- **Plan-as-program**: This view suggests that plans should be interpreted as domain-independent execution directives, which directly determine their user's actions. The users of the plan "blindly" execute the plan with no respect for the current state of the domain.
- **Plan-as-resource** also known as the plan-as-communication view: This view gives plans a smaller role. The plan's users are encouraged to improvise and intelligently perform the tasks suggested by the plan. The user of the plan introduces its own interpretation of the plan depending on the situation.

The situation-dependent interpretation of plans, is performed by the reactive components and sometimes a mediator intermediate level as in ATLANTIS [17], the 3T architecture [9] and Firby's RAPs approach [16]. Maes' view is that there is no need for explicit plan representation and knowledge is directly encoded in the links connecting the behaviors into a network. The architects of DAMN have followed the suggestions of Agre and Chapman word by word by advocating a non-hierarchical approach, where plans are used as sources of information as any other piece of information. Further, the contributions of each behavior to the final decision in DAMN is reflected by a weighting factor generated by the mode manger based on the current context. It is, however, unclear how this mechanism works.

Modifying behavior activation level using weights is used not only in DAMN but in many of the other approaches. In the fuzzy approaches, e.g., [57, 70, 66] another layer of rules determine the activation of the behaviors. In the dynamical systems approach, a competitive dynamic was used to calculate the weights. A case-based approach was used to determine the gain values for motor schemas. And so there are as many approaches to determination of the activation level as there are action selection mechanisms. One thing, however, is obvious: a more global contextual information, than is readily available by reactive behaviors, is used to activate behaviors. We call this approach *context-dependent activation of behaviors.* The general idea is that the behaviors of the system are activated according to the context in which the robot is. These techniques require robust ways of determining the current context.

13.8 Hierarchy vs. no hierarchy

By contrasting the none hierarchical approach utilized in DAMN with the hierarchical approaches in, e.g., ATLANTIS and similar systems, the following question raises: *should hybrid systems be hierarchical or non-hierarchical?*

The logical architecture of most hybrid systems is layered according to decreasing levels of "deliberation" or "consciousness" from top to bottom. The level of abstraction also decreases from top to bottom, with high-level planning components operating on symbolic descriptions and the low-level reactive behaviors operating at the signal level. Abstraction may be an advantage when modeling complex systems. When going from one level to the other information is usually translated to other forms (from signal to symbol and vice-versa). This interaction can complicate the interface between the layers and introduce bottle neck effects.

As a contrast in DAMN all system components, planning or reactive, communicate using a specified language of preferences. There is no information flow between the modules, only the voter collects and processes the preferences and then combines them. This reduces the amount of interaction going on and actually allows for a simpler interface which is homogeneous for all types of system components. Further, reduced interaction increases system modularity which is very attractive from a software engineering point of view. Principally, in this approach all the modules operate at all times in order to vote for the action that they prefer. While this might work for a small number of modules, it will run out of resources with many modules. Hierarchical systems focus the agent's attentions (and hence use of its resources) on a specific task gradually and from top to bottom. For example a step in a plan only requires a subset of the system modules to be invoked which in turn involves a subset of the reactive modules at the lowest level.

As mentioned above, one important characteristic of hierarchical systems is that a decision is made at every level and only a subset of system components are enabled, which means efficient use of system resources. The mechanism used in this kind of hierarchical decision structures is called *rigid switching mechanism*. *Free-flow* hierarchies [69] on the other hand are not governed by a switching mechanism at each level, but rather they place restriction on the flow of information through the hierarchy. Such an architecture would be one similar to DAMN, where the system consisted of a hierarchy of systems and subsystems. Each level would vote for activation of the subsystems immediately below and the lowest level would vote for primitive actions.

13.9 Synchronous vs. asynchronous

In ATLANTIS the view is that interaction between traditional planner and reactive system components should be asynchronous. This allows the fast reactive system components to react in real-time whereas the slow planning components can run at a slower paste without affecting the systems fast response to run-time contingencies.

In this way the reactive components are able to attend to events as they unfold and trigger appropriate responses in a timely manner, since reactive system components have little or no processing and introduce a tight sensor-to-actuator coupling. Asynchronous operation in DAMN is allowed by letting each system component run at whatever speed they can/want.

In CIRCA [39] a component is included that enables reasoning about time consumption of modules and uses this knowledge to trade-off between reactivity and thoughtfulness. I.e., if the system is in no hurry it will take its time and act appropriately and if it is in a hurry it will act faster but less appropriate. In [38] a taxonomic overview of integration of AI and real-time systems is provided.

13.10 Uncertainty handling and reliability

A related problem to the action selection problem is the reliability of the decisions the agent makes.

Reliability in robotics has been studied along several main paths: *reactivity*, *error recovery* and *uncertainty handling*[5]. Reactivity [10, 5] provides immediate responses to unpredictable environmental changes through a tight coupling of perception and action. Reactive architectures have shown improved reliability when compared with classical sense-plan-act architectures. In error recovery techniques a set of dedicated modules monitor the task continuously and upon detection of an error (cognizant failure) an appropriate corrective action is invoked, which can handle unanticipated situations [61, 17]. Further the method utilized in the RAPs approach can be considered as a method exploiting redundancy. Basically in this approach a set of (functionally redundant) modules are invoked until one succeeds. By repeating this process the chances of success are increased. The basic idea behind uncertainty handling techniques is to have explicit "models" of the robot's sensing and action capabilities. Using this model the agent (robot) can predict what actions to take in order to increase the expected utility or the reliability of its task [25, 42]. Even various combinations of ,e.g., reactivity and uncertainty handling have been exploited to an extent. The usual approach is to combine (high-level) deliberative planners that have uncertainty handling capabilities with (low-level) reactive modules (behaviors) which handle run-time contingencies [57, 25].

References

[1] Philip E. Agre and David Chapman. What Are Plans for? *Robotics and Autonomous Systems*, 6:17–34, 1990.

[5] Uncertainty handling can be further divided into the two related areas of *explicit* uncertainty handling and methods exploiting *redundancy* (such as sensor fusion).

[2] Philip E. Agre et al. Markers Elucidated and Applied in Local 3-Space. In *IEEE International Symposium on Computer Vision*, November 1995.

[3] Tracy L. Anderson and Max Donath. Animal Behavior as a Paradigm for Developing Robot Autonomy. *Robotics and Autonomous Systems*, 6:145–168, June 1990.

[4] Ronald C. Arkin. Motor schema based navigation for a mobile robot: An approach to programming by behavior. In *IEEE International Conference on Robotics and Automation*, pages 264–271. IEEE, 1987.

[5] Ronald C. Arkin. Integrating Behavioral, Perceptual and World Knowledge in Reactive Navigation. *Robotics and Autonomous Systems*, 6:105–122, June 1990.

[6] Ronald C. Arkin and Douglas MacKenzie. Temporal Coordination of Perceptual Algorithms for Mobile Robot Navigation. *IEEE Transactions on Robotics and Automation*, 10(3):276–286, June 1994.

[7] Tucker Balch, Gary Boone, et al. IO, GANYMEDE and CALLISTO, A Multiagent Robot Trash-Collecting Team. *AI Magazine*, 16(2):39–51, summer 1995.

[8] Randall D. Beer, Hillel J. Chiel, and Leon S. Sterling. A Biological Perspective on Autonomous Agent Design. *Robotics and Autonomous Systems*, 6:169–186, June 1990.

[9] Pete Bonnaso, Erann Gat, James Firby, and Kortenkamp. A three tire architecture for. Technical report, Johnson Space Center, NASA, 1996.

[10] Rodney A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[11] Rodney A. Brooks. The Behavior Language; User's Guide. *MIT Memo 1227*, April 1990.

[12] Vira Chankong and Yacov Y. Haimes. *Multiobjective Decision Making - Theory and Methodology*, volume 8. North-Holland, 1983.

[13] Christopher I. Connolly and Roderic A. Grupen. On the Application of Harmonic Functions to Robotics. *Journal of Robotic Systems*, 10(7):pp 931–946, October 1993.

[14] Andrew H. Fagg, David Lotspeich, Joel Hoff, and George A. Bekey. Rapid Reinforcement Learning for Reactive Control Policy Design in Autonomous Robots. *Artificial Life in Robotics (T. Shibata & T. Fukuda, eds.)*, 1996.

[15] Cynthia Ferrell. Global Behavior via Cooperative Local Control. *Autonomous Robots*, 2(2):105–125, 1995.

[16] James Firby. *Adaptive Execution in Dynamic Domains*. PhD thesis, Yale University, 1989.

[17] Erann Gat and Greg Dorais. Robot Navigation by Conditional Sequencing. *IEEE Int. Conf. on Robotics and Automation*, 2:1293–1299, 1994.

[18] Joel Hoff and George Bekey. An Architecture for Behavior Coordination Learning. In *IEEE International Conference on Neural Networks, Perth, Australia*, volume 5, pages 2375–2380, November 1995.

[19] Mark Humphrys. *Action Selection methods using Reinforcement Learning*. PhD thesis, Computer Laboratory, University of Cambridge, 1997.

[20] Peter G. W. Keen. The Evolving Concept of Optimality. In Martin K. Starr and Milan Zeleny, editors, *Multiple Criteria Decision Making*, volume 6 of *Studies in the management sceiences*, pages 31–57. North-Holland Publishing Company, 1977.

[21] Omar Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *The International Journal of Robotics Research*, 5(1):90 – 98, 1986.

[22] Kurt Konolige. *Saphira Software Manual*. ActivMedia Inc., version 6.1 edition, 1998.

[23] Jana Koŝecká and Ruzena Bajcsy. Discrete Event Systems for Autonomous Mobile Agents. . *Proceedings Intelligent Robotic Systems '93 Zakopane*, pages 21–31, July 1993.

[24] Jana Koŝecká, Henrik I. Christensen, and Ruzena Bajcsy. Experiments in behavior composition. *Robotics and Autonomous Systems*, 19:287–298, 1997.

[25] Steen Kristensen. Sensor planning with Bayesian decision theory. In *Reasoning with Uncertainty in Robotics*. University of Amsterdam, The Netherlands, December 1995.

[26] Steen Kristensen. *Sensor Planning with Bayesian Decision Analysis*. PhD thesis, Department of Medical Informatics and Image Analysis, Aalborg University, 1996.

[27] Steen Kristensen. Sensor planning with Bayesian decision theory. *Robotics and Autonomous Systems*, 19(3-4):273 – 286, March 1997.

[28] Edward W. Large, Henrik I. Christensen, and Ruzena Bajcsy. Scaling the Dynamic Approach to Path Planning and Control: Competition among Behavioral Constraints. *International Journal of Robotics Research*, 18:37–58, January 1997.

[29] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.

[30] Long-Ji Lin. Scaling up Reinforcement Learning for robot control. In *Proceedings of the Tenth International Conference on Machine Learning*, 1993.

[31] Demian M. Lyons. Representing and Analyzing Action Plans as Networks of Concurrent Processes. *IEEE Transactions on Robotics and Automation*, 9(7):241–256, June 1993.

[32] Douglas C. Mackenzie, Ronald C. Arkin, and Jonathan M. Cameron. Specification and Execution of Multiagent Missions. *Autonomous Robots*, 4(1), Januray 1997.

[33] Pattie Maes. How To Do the Right Thing. Technical Report NE 43 - 836, AI-Laboratory, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139, USA, 1989.

[34] Pattie Maes. Situated agents can have goals. In Pattie Maes, editor, *Designing Autonomous Agents*, pages 49–70. MIT/Elsevier, The MIT Press, Cambridge, Massachusetts, London, England, 1990.

[35] Sridhar Mahadevan and Jonathan Connel. Automatic programming of behavior-based robots using reinforcement learning. *Artificial Intelligence, Elsevier*, 55:311–365, 1992.

[36] Maja Mataric. Behavior-Based Control: Examples from Navigation, Learning, and Group Behavior. *Journal of Experimental and Theoretical Artificial Intelligence, special issue on Software Architectures for Physical Agents*, 9(2-3):323–336, 1997.

[37] Marvin Minsky. *The Society of Mind.* Simon and Schuster, 1985.

[38] David J. Musliner, James A. Hendler, Ashok K. Agrawala, Edmung H. Durfee, et al. The Challenges of Real-Time AI. *IEEE Computer*, 28(1), January 1995.

[39] D.J. Musliner, E.H. Durfee, and K.G. Shin. CIRCA: A Cooperative Intelligent Real-Time Control Architecture. *IEEE Transaction on Systems, Man, and Cybernetics*, 23(6):1561–1574, November/December 1993.

[40] A. Newell and H. Simon. *GPS: A program that simulates human thought.* Computers and Thought. McGraw-Hill, 1963.

[41] Illah Nourbakhsh, S. Morse, et al. The Winning Robots from the 1993 Robot Competition. *AI Magazine*, 4(14):51–66, Winter 1993.

[42] Illah Nourbakhsh, Rob Powers, et al. DERVISH, An Office-Navigating Robot. *AI Magazine*, 16(2):53–60, summer 1995.

[43] David W. Payton. Internalized plans: A representation for action resources. In Pattie Maes, editor, *Designing Autonomous Agents*, pages 89–104. MIT/Elsevier, The MIT Press, Cambridge, Massachusetts, London, England, 1990.

[44] David W. Payton, J. Kenneth Rosenblatt, and David M. Keirsey. Plan guided reaction. *IEEE transactions on Systems, Man, and Cybernetics*, 20(6):1370–1382, Novermber/December 1990.

[45] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* The Morgan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann, San Mateo, California, 1988.

[46] Paolo Pirjanian. *Multiple Objective Action Selection & Behavior Fusion using Voting*. PhD thesis, Department of Medical Informatics and Image Analysis, Institute of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7, DK-9220 Aalborg, Denmark, August 1998. *Available online: http://www-robotics.usc.edu/˜paolo/publications.*

[47] Paolo Pirjanian and Hans Blaasvær. AMOR: An Autonomous Mobile Robot Navigation System. Master's Thesis, Laboratory of Image Analysis, Institute of Electronic Systems, Aalborg University, Denmark, 6 1994.

[48] Paolo Pirjanian and Maja Matarić. A decision-theoretic approach to fuzzy behavior coordination. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, November 1999.

[49] Paolo Pirjanian and Maja J Matarić. *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*, chapter Multiple Objective vs. Fuzzy Behavior Coordination. Lecture Notes in Computer Science, 2000. To appear.

[50] Paolo Pirjanian and Maja J Matarić. Multi-robot target acquisition using multiple objective behavior coordination. In *IEEE Conference on Robotics and Automation*, San Francisco, April 2000. To appear.

[51] Ashwin Ram, Ronald C. Arkin, et al. Case-Based Reactive Navigation: A case-based method for on-line selection and adaption of reactive control parameters in autonomous robotic systems. Technical Report GIT-CC-92/57, College of Computing, Georiga Institute of Technology, Atlanta, Georgia 30332-0280, USA, 1992.

[52] Jukka Riekki and Juha Röning. Reactive Task Execution by Combining Action Maps. In *International Conference on Integrated Robots and Systems (IROS)*, pages 224 – 230, Grenoble, France, September 1997.

[53] Julio K. Rosenblatt. DAMN: A Distributed Architecture for Mobile Navigation. In *AAAI Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents, Stanford, CA.* AAAI Press, Menlo Park, CA., March 1995.

[54] Julio K. Rosenblatt and Charles E. Thorpe. Combining multiple goals in a behavior-based architecture. In *International Conference on Integrated Robots and Systems (IROS)*, August 1995.

[55] Stuart Russell and Eric Wefald. *Do the Right Thing - Studies in Limited Rationality*. The MIT Press, Cambridge, Massachusetts, London, England, 1991.

[56] Alessandro Saffiotti. The Uses of Fuzzy Logic in Autonomous Robot Navigation: a catalogue raisonné. Technical Report 2.1, IRIDIA, Universite Libré de Bruxelles, 50 av. F. Rossevelt, CP 194/6, B-1050 Brussels, Belgium, November 1997.

[57] Alessandro Saffiotti, Kurt Konolige, and Enrique H. Ruspini. A multivalued logic approach to integrating planning and control. *Artificial Intelligence*, 76:481–526, March 1995.

[58] Alessandro Saffiotti, Enrique Ruspini, and Kurt Konolige. *Handbook of Fuzzy Sets and Possibility Theory*, chapter Using Fuzzy Logic for Mobile Robot Control. Kluwer Academic, 1997.

[59] Gregor Schöner and M. Dose. A dynamics systems approach to task level systems integration used to plan and control autonomous vehicle motion. *Robotics and Autonomous Systems*, 10:253–267, October 1992.

[60] Schopenhauer. *Schopenhauers Filosofi (eng. Schopenhauer's Philosophy)*. Nyt Nordisk Forlag, 1947.

[61] Reid G. Simmons. Structured Control for Autonomous Robots. *IEEE Trans. on Robotics and Automation*, 10(1):34–43, February 1994.

[62] Herbert A. Simon. *The New Science of Management Decision*. Harper and Row, New York, 1960.

[63] Martin K. Starr and Milan Zeleny. MCDM-State and Future of the Arts. In Martin K. Starr and Milan Zeleny, editors, *Multiple Criteria Decision Making*, volume 6 of *Studies in the management sceiences*, pages 5–30. North-Holland Publishing Company, 1977.

[64] Axel Steinhage and Gergor Schöner. The Dynamic Approach to Autonomous Robot Navigation. In *IEEE International Symposium on Industrial Electronics*, pages 7–12, Piscataway, NJ, USA, 1997.

[65] John K. Tsotsos. Behaviorist Intelligence and the Scaling Problem. *Artificial Intelligence*, 75(2):135–160, June 1995.

[66] Eddie Tunstel. Coordination of Distributed Fuzzy Behaviors in Mobile Robot Control. In *IEEE International Conference on Systems, Man,and Cybernetics*, pages 4009–4014, Vancouver, BC, Canada, October 1995.

[67] Roy M. Turner. Context-sensitive Reasoning for Autonomous Agents and Cooperative Distributed Problem Solving. http://bronte.umcs.maine.edu/Papers/1993/ijcai-ctx-rmt/tr.html, University of New Hampshire, 1996.

[68] Roy M. Turner, D. Richard Bildberg, et al. Generic Behaviors: An Approach to Modularity in Intelligent Systems Control. http://cdps.umcs.maine.edu/Papers/1993/auv-rmt/tr.html, 1993.

[69] Toby Tyrrell. *Computational Mechanisms for Action Selection*. PhD thesis, University of Edinburgh, 1993.

[70] John Yen and Nathan Pfluger. A Fuzzy Logic Based Extension to Payton and Rosenblatt's Command Fusion Method for Mobile Robot Navigation. *IEEE Transactions on Systems, Man, and Cybernetics*, 25(6):971 – 978, June 1995.