

Integrating Effective Planning Horizons into an Intelligent Systems Architecture

J. P. Gunderson and L. F. Gunderson
Gunderson and Gunderson, Inc.

ABSTRACT

One metric of the intelligence of a system is its ability to perform tasks in the face of dynamic changes to the environment. This requires that an autonomous system be capable of responding appropriately to such changes. One such response is to effectively adapt the allocation of resources from planning to execution.

By adapting the resource allocation between cognition and execution, an intelligent system can produce shorter plans more frequently in environments with high levels of uncertainty, while producing longer, more complex plans when the environment offers the opportunity to successfully execute complex plans.

The effective planning horizon is developed from an analysis of mathematical models of classic autonomous system and from current research in cognitive science. Experimental results are presented showing the performance gain from an effective planning horizon based system.

From this simplified feedback control model, the Effective Planning Horizon concept is extended to a more realistic intelligent system architecture, and the concepts of bounded rationality, intelligent heuristics and the judgment analysis "lens model" are shown to be analogs of the effective planning horizon.

Keywords: *Effective Planning Horizon, Interleaved Plan/Execution, Probabilistic Planning, Intelligent System.*

1 INTRODUCTION

The world is not a perfect place, and if intelligent systems are to function effectively they must be capable of handling both the uncertainties of the world outside themselves and those created by their own limitations. One working definition of intelligence is the ability to adapt to unexpected failure. Using this definition, intelligence is defined by the ability of the system to achieve goals in the face of failure. In order to do this, the system must have some mechanism to detect that its current method for satisfying its goals is insufficient, and to select and apply new methods.

In this paper we present the concept of an effective planning horizon as a requirement for intelligence. We start by discussing the reasons that an effective planning horizon is necessary. We place the model of an effective planning horizon into a model of an intelligent system. Then we develop a mathematical definition of the effective planning horizon. This definition is tested in the domain of a simulated Autonomous Underwater

Vehicle. In the final sections of the paper we present arguments to show that the effective planning horizon concept can easily be transformed into its analogs in sensing, acting, and goal selection.

2 WHY AN EFFECTIVE PLANNING HORIZON

Why can't an intelligent system just make one plan? The answer to this lies in the balance between the complexity of all environments and the limited computational resources that any intelligent system can apply to the problem.

It is obvious that all natural environments and all but the simplest of synthetic environments are stochastic. Even in a simple deterministic environment, the observations made by an entity have a probability distribution [1]. As demonstrated below, if the environment is not deterministic, then the longer the plan needed to meet the goal, the less likely it is to succeed. However, even in an interleaved planning-execution system, too short a plan length will often lead to sub-optimal plans.

Alternatively a system could generate all conceivable plans in advance and maintain those in a plan library [2]. But, in any realistic domain, it is not possible for an intelligent system to completely explore all the possible ramifications of a plan of action in a realistic time [3]. In addition, to time constraints, there are memory constraints. In fact for one nominally intelligent system, only 7 ± 2 objects can be held in working memory at any one time [4]. This seriously limits the number of plans that can be considered. This has led to the concepts of "bounded rationality" in both biologic [5] and machine based intelligence [6].

Thus, it appears that all intelligent systems (whether biologic or electronic) must make trade-offs between cognition and action. There are many mechanisms for placing bounds on cognition, and in this paper we focus on the concept of a planning horizon.

3 A MODEL OF AN INTELLIGENT SYSTEM

The model we use is based on the elementary loop of functioning (E.L.F.) model of Meystel and Albus [6]. This model assigns responsibility for four critical tasks to independent processes in an iterative - *sense, model,*

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE AUG 2002		2. REPORT TYPE		3. DATES COVERED 00-00-2002 to 00-00-2002	
4. TITLE AND SUBTITLE Integrating Effective Planning Horizons into an Intelligent Systems Architecture				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Gunderson and Gunderson, Inc., Gamma Two Inc, 209 Kalamath St, Denver, Co, 80223				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Proceedings of the 2002 Performance Metrics for Intelligent Systems Workshop (PerMIS ?02), Gaithersburg, MD on August 13-15, 2002					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

plan, act loop. We have extended this model by adding a new process, *validate*. This new process has the ability to compare the observed result of the previous cycle with the expected result. It then updates the other processes with the observed success or failure of the cycle. This is shown graphically in Figure 1.

This model allows for the representation of intelligence as a semi-lattice. In this representation, the intelligence of a system is defined by the ability of the system to perform each of these processes. We utilize goal satisfaction rates as a primary metric of intelligence. The intelligence of the system is a 5-tuple of the intelligence of each of the processes. The intelligence of each of the processes is incommensurate with the intelligence of the other processes, thus establishing the partially ordered set needed for the lattice.

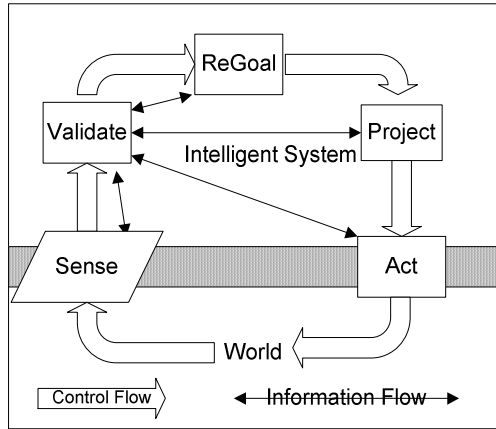


Figure 1 Extended model of Intelligent Systems

4 INTELLIGENT SYSTEMS

Since the earliest work in machine intelligence, a concern has been the computational burden required for ‘intelligent’ behavior. With the representation of planning as search [8] came the exponential growth of the number of *world states* explored with the increase in search depth. The number of world states explored is a measure of the computational complexity of the planning process. It has been long accepted that as uncertainty and dynamism increase in an environment, the need for more reactive systems also becomes greater [9]. In recent work, quantitative relationships have been suggested for the impact of three types of domain uncertainty on the ability of an autonomous system to achieve goals.

This work classifies domain uncertainty into three categories:

- Sensor uncertainty;
- Effector uncertainty; and,
- Uncertainty caused by exogenous events.

These three types of uncertainty have been used in a simulation system that measures the goal satisfaction of a simulated maintenance robot under widely varying levels of each of these types of uncertainty [10].

These results suggest that autonomous systems are very sensitive to even low levels of uncertainty in the environment, with overall goal satisfaction dropping to 75% with the introduction of 10% sensor errors. In addition, this research discovered that the ability to retry on failure was very effective at allowing the system to maintain goal satisfaction in the face of all types of uncertainty. The simulated robot was more robust (i.e., able to satisfy goals as the environment changed) when it used different deliberation and execution mixes in response to different levels of uncertainty in the environment.

5 DEFINING THE EFFECTIVE PLANNING HORIZON

Beginning with early two-player game models, it has been common to place a limit on the search depth used to explore options. This was often implemented as an arbitrary limit on the ply-depth of an alpha-beta search tree.

In two player games of perfect information, it is implicitly assumed that if the computational resources were sufficient, the ply-depth could be increased until a perfect game analysis was performed. However, when an intelligent system is playing against the real world, conditions change. Since there is an inherent variability in the world, perfect information is impossible. Under these conditions, there appears to be a hard limit on the planning horizon – a line beyond which planning is useless, since it is impossible to know the state of the world in which the planning is taking place.

We approach the idea of an effective planning horizon from the perspective of the expected value of additional planning. What is the value associated with extending my plan, and what is the likelihood that the world will be in the state I planned for?

$$E(plan) = Pr(plan) * Value(Plan) \quad (1)$$

While solving the exact value of equation 1 may be impossible, it is possible to make some reasonable assumptions about the terms, and from these assumptions draw some conclusions about the characteristics of the expected value of planning in dynamic and uncertain domains.

We make three assumptions about the actions in a plan:

1. The action achieves some state that is needed to achieve a goal (or goals);

2. Each action has some finite probability of failure; and,
3. All the actions in the plan must succeed for the plan to succeed.

Under these assumptions it is clear that a first approximation of the probability of the plan succeeding is:

$$\Pr(plan) = \prod_{steps} \Pr(Step) \quad (2)$$

This can be further simplified by assuming that all actions have some common probability of success, reducing Equation 2 to:

$$\Pr(plan) = \Pr(step)^{length} \quad (3)$$

Approximating the value associated with a plan of given length is more problematic, however for any given domain and set of possible goals to achieve some useful assumptions can be made.

First, let us envision the complete range of problems that our intelligent system might be required to solve. This range could be defined as the cross product of all possible initial conditions with all possible final states. Then let us imagine the plans that might be utilized to transform the world as we find it into the world as we desire it. Clearly, very few of the (Initial, Final) pairs can be transformed by plans utilizing a single action. A few more, perhaps, can be achieved with two-action plans, more yet by three action plans, and so on. If we further take the view that we wish the simplest, or most likely, plan to satisfy our needs, then there will be few pairs that require 1000 action plans, and fewer still that require 2000 action plans. If we model the value of a plan of size N as the number of (Initial, Final) pairs that can be satisfied by a plan of length N or less, we get a sigmoidal, or logistic curve. This curve can be closely approximated with an equation like (3) below:

$$Value(plan) = 1 - \frac{a}{1 + 2^{length}} \quad (4)$$

Finally, combining the two terms into the expected value of a plan of length N, for a given domain, and task assignment, we find that the resulting curve has the characteristic shape shown in Figure 2.

This finding is consistent with the analysis that, there exists some point beyond which the costs associated with planning exceed the benefits, which causes the expected value to decrease. Under the assumptions made above, we can conclude that there is a range of plans, which provide the optimal benefit to the deliberative

system. This range is defined as the effective planning horizon (EPH). If the planning depth is less than the EPH, the probability that effective solutions will be produced is too low. If the planning depth exceeds the EPH, the probability that significant amounts of computational resources will be expended planning for situations that never occur is high. If those resources had been applied more effectively, the rate of goal satisfaction could increase. However, while this analysis suggests that it is beneficial to adjust the planning horizon to the domain and goals of the intelligent system, it does not provide any mechanism to achieve this adjustment.




Figure 2 - Expected Value of Planning

5.1 METHODS OF ADJUSTING THE PLANNING HORIZON

In this paper we investigate a method of selecting an effective planning horizon, adapted directly from the notion of placing a limit on the search depth. The planning system used is an interleaved, probability-aware, forward chaining planning and execution environment. It is a general-purpose planner in that all domain specific information is encoded as part of an input file. Included in this domain specific information are naive probabilities of success for the available actions. The planning system uses this information to produce the plan with the highest observed probability of success.

These experiments were run using an autonomous underwater vehicle domain. In this domain the planner is embedded in the autonomous underwater vehicle (AUV). The types of tasks assigned to the AUV include autonomous navigation, the deployment of monitoring devices within enemy controlled territory, and avoiding or escaping detection by enemy Anti-Submarine Warfare (ASW) surface vessels.

A typical task assignment would be to:

1. Begin from Home base, at the surface, carrying a deployable monitor;

2. End at Home base, at the surface, with the monitor deployed at a location just outside the mouth of the enemy harbor; and,
3. By the way, don't get detected by any enemy ships while you do it.

If the system has perfect knowledge, and nothing changes, this task is a straight forward planning exercise. However, in this domain, the enemy ship moves while the AUV is executing its mission and the AUV has limited sensor range. This forces the planner to develop a plan, begin to execute it, and respond to failures. These failures could include finding the ASW ship in its path or being detected during a move. More complex failures might involve having to switch goals to escape detection and then re-target the goal of dropping the monitor, once it has lost its opponent.

The planning and execution system is interleaved, so during the execution of actions, it receives feedback about the actual state of the world (at least those parts which it can sense) and it compares the actual state with its expected state. As long as the expected state of the world and the actual state agree, the system continues to execute the planned actions. If the states do not match, several options are available:

1. Continue with the current plan,
2. Develop a new plan to meet the current goals,
3. Re-evaluate the goals, and develop a plan to meet the new goals, or
4. Do nothing, and hope that the world changes.

After this evaluation, a new action is issued, and the process begins again. Since the planner is deeply embedded in the execution loop, strict limits on computation must be imposed to assure responsive behavior of the system.

5.2 SEARCH DEPTH CONTROL

Search depth control has been used to limit computation for as long as planning as search has been used. In general there are two forms of this control: limiting the exploration depth directly – such as ply-depth limits, or plan length limits- and world set limits. In the latter case, the planning system monitors the total number of individual worlds explored, and at some pre-determined limit, stops exploration. This has benefits when coupled with search control rules which allow the planner to ‘focus attention’ on areas of the plan which might be more fruitful. The experimental planning system used here uses this world count form of search depth control.

To explore the impact of limiting search depth on plan success rates the scenario is analyzed under a range of allowable search depths.

6 EXPERIMENTAL RESULTS

The experiments were designed to answer the question “For a given domain, and a given task mix, at what point

does it stop paying to plan?” Using the analysis in Section 5, we measure the marginal value of planning by measuring the success rate of the plans produced, and the computational complexity of producing the plan. All experiments are run on the same interleaved planning and execution system, coupled with an external simulator of the domain.

6.1 EXPERIMENT SETUP

The domain used is that of the AUV, described above. The AUV can navigate in a world of 18 possible locations, at up to three depths for each location. Any motion has a risk of causing the AUV to be detected, however the deeper the AUV the lower the risk, In addition to movement actions, the AUV can change depth, deploy a monitor, and cause a tracking ASW vessel to lose track by going deep and drifting. The AUV carries one monitor, and the task requires the AUV to deploy that monitor at a specific location and return to base undetected.




Figure 3 - Autonomous Underwater Vehicle domain

The task assigned to the AUV in these experiments is to travel from home base (Location A) to Location P, release a monitor, and return to Location L for pickup (sub-lp10.pro). It has the additional requirement of remaining undetected during the mission, and should it be detected it must cause the enemy vessel to lose its lock before being picked up (Figure 3). The AUV has six general types of operators to apply:

Table 1 Relative probability of success for the AUV domain. Failure can mean either the AUV is detected, the operator fails, or both.

Operator Type	<i>A priori</i> Probability of Success
---------------	--

Travel Surface	0.45
Travel Shallow	0.90
Travel Deep	0.95
Change Depth	0.99
Release Monitor	0.99
Drift (can cause opponent to lose detection)	0.25

For the assigned task, the naïve probability of success is approximately 0.46, if the plan with highest probability is selected, or 0.00015 if the plan with the minimum number of operators is selected. One additional complicating factor is the behavior of the ASW vessel. This unit moves around the map in a random walk, and if it occupies a location it blocks the AUV from traveling to that location. The location of the ASW is initially unknown by the AUV, however, the AUV can sense the presence of the enemy vessel if it enters or leaves an adjacent location.

A single simulation consists of assigning a task, executing the mission, and recording the number of individual world states explored by the planner during the execution of the mission and the success or failure of the mission.

Two baseline simulations are initially run. In the first the planning horizon required for success by a non-interleaved planning system is determined. Since a non-interleaved system must be able to produce a single plan to meet all the mission goals, the required search depth is much greater than that required by an interleaved planner. In this experiment the planning horizon is increased until the planning system reaches the point where it can produce a complete, successful plan. This established the minimum planning horizon for a non-interleaved planner. Since the planner is only allowed to execute a single cycle, the success rates of this plan cannot be effectively compared with the success rates of the interleaved simulations.

The second baseline addresses this problem by using the planning horizon established by the first baseline. However, the planning/execution system is run in interleaved mode, allowing the planner to correct failures. This second baseline establishes the plan success rate and the minimum number of world states required for an interleaved planner operating without a search depth limit.

The experimental simulations use an interleaved planning system, with a range of planning horizons. In all experiments a minimum of 100 independent simulations are run at each planning horizon. In interleave mode, the system is allowed to compete up to five planning/execution cycles. The average number of world states examined during the each simulation and the average success rate of the mission are recorded.

6.2 DATA ANALYSIS

The data analysis is straightforward. We calculate the value of the planning horizon (PH) as:

$$Value(PH) = \frac{SuccessRate}{MegaWorldsExplored} \quad (5)$$

The value is the success rate over the work expended (million worlds explored to achieve the success rate).

Baseline 1 established that it is necessary to examine approximately 13,000 world states, and the success probability is 0.14 for the non-interleaved planning system. Due to the interaction with the ASW vessel, this is significantly below the naïve probability (0.46). In effect, most of the computational resources used by the intelligent system are expended planning for situations that simply do not occur during the execution. For Baseline 1 the Value(PH) = 10.76.

Using the planning horizon established by Baseline 1, Baseline 2 achieves a success ratio of 0.74. However it explores 66,000 world states, for a Value(PH) = 11.21. Interestingly, even though the success ratio increased, the number of world states explored did so proportionally. This data is presented in Table 2.

Table 2 Baseline performance of complete planning system. In Baseline1 a single iteration of the planning system was allowed, in Baseline2 five iterations were completed.

Test	Worlds	Success Rate	Value(PH)
Baseline1	13,000	0.14	10.76
Baseline2	66,000	0.74	11.21

The experimental simulation was run with the range of planning horizons shown in Table 3. Note that the peak value of value(PH), 68.7 is approximately 6 times the value achieved by Baseline 2. This shows that at the effective planning horizon, the planning system is achieving equivalent success rates with significantly less computational cost.

Table 3 Table of the computational resources required for varying planning horizons, the achieved success rates, and the Value of the invested computational resources.

Planning Horizon	Worlds	Success Rate	Value(PH)
100	9,605	0.02	2.08
150	14,200	0.04	2.81
200	18,854	0.06	3.18
250	12,813	0.88	68.7

300	14,586	0.88	60.3
350	16,717	0.91	54.4
400	19,358	0.89	46.0
450	20,026	0.94	46.9
500	22,315	0.90	40.3
550	22,726	0.93	40.9
600	27,615	0.86	31.1

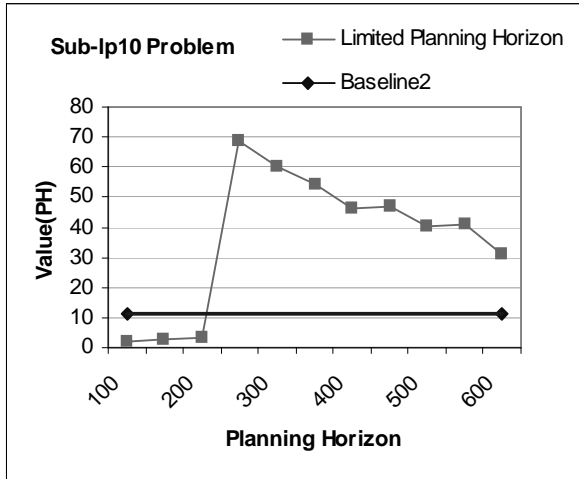


Figure 4 - Plotted Value of Effective Planning Horizons from the AUV problem

In Figure 4, the value of the planning horizon, as defined in Equation 5 is plotted against the search limit imposed on the planner. The first three points are characteristic of planning horizons that are too short to allow the planning system to find successful plans. The range from 250 – 350 represents the range of the EPH where successful plans are being found with minimum computational investment. The range above 350 is characteristic of wasted resources, planning for situations that never occur.

The effective planning horizon is the global maximum of this curve. It should be noted, that this figure closely resembles the analytic solution shown in Figure 2. The “Baseline2” line is the value achieved by Baseline 2.

It is clear that limiting the planning horizon in interleaved intelligent systems allows the system to achieve success rates which are comparable to complete planning systems, at a fraction of the invested resources. While the current data results from adjusting the search limit utilized by the planning system, several other mechanisms exist to adjust the planning horizon, including waypoint based planning [11], and the use of assumptive systems [12], which was successfully applied during the 1994 AAI Robot Contest.

7 EXTENSION TO A COMPLETE INTELLIGENT SYSTEM MODEL

The preceding discussion describes the benefits of an effective horizon in planning. However, this argument can be extended to all of the other processes.

- Sensing
- Validation
- Re-Goaling
- Acting

In the next section we sketch some ways in which effective horizons might be achieved for these processes. Much of this discussion will be derived from work done on biologic systems.

7.1 EFFECTIVE SENSING HORIZONS

One model of the sensing process can be taken from the work of Egon Brunswik [1]. In this lens model, shown graphically in Figure 5, perception can be modeled as a linear weighted sum:

$$y_s = \sum_{i=1}^n w_i x_i \quad (6)$$

- Where y_s = the judgment of the condition of target s
- y_e = the actual environmental condition of the target
- n = the total number of cues available to the judgment maker
- x_i = value of cue i , where i goes from 1 to n
- w_i = the weighting of cue i

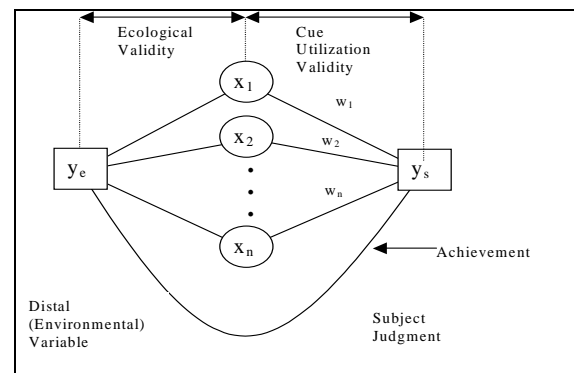


Figure 5 - The Lens Model of Egon Brunswik

For this process the effective sensing horizon will be a balance between using so few cues that the perception is invalid and using so many cues that the cognitive system is overloaded.

7.2 EFFECTIVE VALIDATING HORIZONS

Validation of the successful completion of a cycle is critical if the intelligent system is to adapt to unexpected

outcomes. However, given limited sensing ranges, and long distance effects, it can be resource intensive to do complete validation. In some cases, the system may have to develop and achieve complete new sub-goals to accomplish this phase. Yet, the costs of proceeding without validation can be extreme. Thus any intelligent system must strike a balance between the costs and benefits of validating the world state.

In addition, when things go wrong, an intelligent system must assess the probable causes of the failure, and make improvements or the same ineffective actions will occur again and again. If the mission failed because a sensor was incorrectly interpreted causing a specific action to fail, it is not intelligent to expend resources improving the action. However, correctly attributing the source of a failure is extremely resource intensive, and limits must be placed on its application.

7.3 EFFECTIVE GOAL RE-PRIORITIZATION HORIZONS

Intelligent systems do not pursue the same goals at all times. Consider the mother alligator, who under normal circumstances has a hair trigger bite reflex. Yet, this same alligator will carefully carry its young about cradled in those same jaws. Clearly, this intelligent system is re-prioritizing its goals in response to changes in its environment.

However, this ability to re-goal comes at a cost. Re-goaling requires the system to use cognitive resources that could be applied in other ways, and errors in goal prioritization can reduce the success probabilities of the system. This is the same dynamic tension that exists in the selection of an EPH.

7.4 EFFECTIVE ACTION HORIZONS

Just as all biologic intelligent systems have limits on cognition, there are limits on both the number and range of actions the system can undertake to achieve a goal, and the quality of those actions. In biologic systems "Use it or lose it" applies, yet limited time is available for practice. With mechanical systems one tends to think of the range of actions that are available as fixed at the time of construction, and the quality (probability of success) as constant. However, as bearings wear out, and rubber gripping fingers age, the ability of the system to meet its goals degrades, until new resources are applied. While self-repair is beyond the current capabilities of most machine-based intelligent systems, they can possess the ability to update the reliability of actions to reflect changes in the system itself.

8 CONCLUSIONS

The focus of this paper is on utilizing limited computational resources to improve the effectiveness of intelligent systems. Drawing on research from existing biological intelligent systems and current research into

machine-based intelligent systems, an analytic definition of the Effective Planning Horizon was developed.

Success probability in stochastic domains was presented as the key metric for evaluating intelligent systems, and several secondary measures including the expected value of planning, and plan value as a function of computational load were introduced as supporting concepts for the Effective Planning Horizon.

The EPH is a measure of the dynamic and variable nature of the environment, and the goals and limitations of the intelligent system. From these characteristics it is possible to establish a horizon beyond which additional planning is ineffective.

A simulated domain was presented which is representative of the types of tasks we expect deployed intelligent systems to undertake, and using a planning and execution system designed for these demanding domains, experimental data was collected. The data collected demonstrates that it is possible to achieve the same levels of success as the computationally expensive complete exploration of a plan space, at significantly lower cost. This lowered cost translates into lowered demands on the system, or increased speed of execution by the intelligent system, which can be crucial design requirements of future embedded intelligent systems.

Finally, the same principles which led to the formalization of the Effective Planning Horizon were applied to the other processes in feedback control theory based intelligent systems, suggesting other mechanisms that can be used to improve goal satisfaction by intelligent systems.

9 REFERENCES

- [1] Brunswik, E. *Perception and the Representative Design of Psychological Experiments*, University of California Press, Berkeley, CA, 1956.
- [2] Schopper, M. J. "Universal Plans for Robots in Unpredictable Environments," Proceedings of the 10th International Joint Conference on Artificial Intelligence, 1987, pp. 1039 – 1046.
- [3] Ginsberg, M. L. "Universal Planning: An (Almost) Universally Bad Idea," *Artificial Intelligence*, vol. 10, num. 4, pp. 40-44, 1989.
- [4] Miller, G. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," *The Psychological Review*, vol. 63, pp. 81-97, 1956.
- [5] Gigerenzer, G. and Todd, P. M. *Simple Heuristics That Make Us Smart*, Oxford: Oxford University Press, 1999, ch. 1, pp. 3-34.
- [6] Bratman, M. E. Israel, D. J. and Pollack, M. E. "Plans and Resource-Bounded Practical

- Reasoning,” *Computational Intelligence*, vol. 4, num. 4, pp. 349-355, 1988
- [7] Meystel, A.M. and Albus, J.S. *Intelligent Systems Architecture, Design, and Control*. N.Y.: Wiley Interscience, 2002, ch. 2, pp 90-92
- [8] Newell, A. and Simon, H. A., *Computers and Thought*, Feigenbaum, E. A. and Feldman, J., eds., N.Y.:McGraw-Hill, 1963,
- [9] Brooks, R. A. “A Robust Layered Control System for a Mobile Robot,” *IEEE Journal of Robotics and Automation*, vol. RA-2(1), pp. 14-23, 1986
- [10] Gunderson, J. P. and Martin, W. N. “The effects of Uncertainty on Plan Success in a Simulated Maintenance Robot Domain,” *Journal of Theoretic and Experimental Artificial Intelligence*, vol. 12, pp153-164, 2000
- [11] Gunderson, J. P. and Martin, W. N. “Applications of Data-Mining to Sub-Plan Selection in Autonomous Planning Systems,” in Proceedings of the IEEE Systems, Man and Cybernetics Conference, 2001, pp 1459-1464.
- [12] Nourbakhsh, I. R., *Interleaving Planning and Execution for Autonomous Robots*, Boston: Kluwer Academic Publishers, 1997, ch. 7, pp 99-107.