Pervasive Reasoning: Integration of Planning and Health Information Gathering

Lukas Kuhn [1]

[1] *PARC, Palo Alto, CA, USA*
Lukas.Kuhn@parc.com

ABSTRACT

In model-based systems, a planner uses a system description to create a plan that achieves operational goals. The same description can be used by a model-based reasoner to infer the condition of components from partially informative sensors. Rasoning makes typically either inferences from passive observations during operation, or takes diagnostic actions while operation must be halted. In *Pervasive Reasoning* the planner exploite plan flexibility to construct plans that simultaneously achieve operational goals while maximizing information gain for reasoning. Our results demonstrate that pervasive reasoning obtains higher long-run productivity than a decoupled combination of planning and reasoning.

1 INTRODUCTION & PROBLEM

Artificial intelligence has long been animated by the vision of creating fully autonomous systems that not only act on the larger world, but also maintain and optimize themselves. The increased autonomy, reliability and flexibility of such systems are important in domains ranging from space craft, to manufacturing processes. Autonomy can be seen as the combination of two processes: reasoning about the current condition of components in a system from weakly informative sensor readings and planning of system operation optimized for the current condition and operational goals. In an aerospace domain, flight dynamics models can be used to reason about faults in flight control surfaces from noisy observations of flight trajectories. A model-based flight planning could then compensate for the faults by using alternative control surfaces or engine thrust to achieve the pilot's goals.

Reasoning and model-based planning are typically combined in one of two ways: 1) alternation of an *explicit reasoning* mode with an operational model-based planning mode or 2) parallel execution of a *passive reasoning* process with an operational model-based planning mode. The explicit reasoning part of the alternation strategy typically results in long periods in which regular operation is suspended in favor of reasoning. This is particularly true when reasoning about rare intermittent events. The combination of a *passive reasoning* process with model-based planning is often unsuccessful as regular operation may not sufficiently exercise the underlying system to gather enough information.

Driven by the goal to enable systems to provide sufficient information for reasoning without suspending regular operation, I got interested in the question: Can we optimize for both operational goals and active information gathering *simultaneously*?

2 PERVASIVE REASONING: CURRENT WORK & PROPOSED PLAN

2.1 Current work: The theoretical concept of *Pervasive Reasoning*

In the first phase of my PhD I investigated together with my advisor and his research group the stated problem in depth and developed a new paradigm, called *Pervasive Reasoning*.

In *Pervasive Reasoning* certain parameters of operational model-based planning are actively manipulated to maximize information gathering while *simultaneously* achieving regular operation. The primary objective in pervasive reasoning is to continue regular operation. Under the assumption that there are various ways to achieve regular operation, pervasive reasoning simultaneously maximizes information gathering. Similar to cost efficient operation, where operational plans are evaluated by cost and the most cost efficient operational plan dominates, pervasive reasoning evaluates the set of plans that achieve operational goals by their potential information gain for reasoning and selects the most promising plan.

The overall goal of pervasive reasoning is to enable a system to accurately determine its state, intelligently reason about its diverse capabilities, and robustly control its behavior by integrating techniques such as task planning, model-adaptation, structure learning and active diagnosis.

2.2 Current work: A Framework for *Pervasive Reasoning*

As part of my PhD I developed a framework around the theoretical concept of *Pervasive Reasoning* depicted schematically in Figure 1. The planner creates an in-

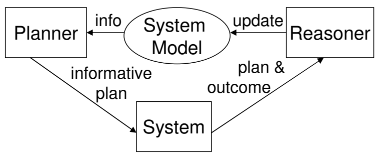

Figure 1: Framework of Pervasive Reasoning for the Optimization of Long Run Efficiency

formative plan that achieves operational goals and at the same time generates potentially informative observations. This plan is then executed on the system and observations are gathered. The reasoning engine updates its beliefs to be consistent with the plan executed by the system and the resulting observations. The reasoning engine forwards updated beliefs to the planner. Based on the updated belief state the planner reasons about which new information is needed by the system. The planner and reasoning engine both run in real time. The framework is built around the concept of maximizing long-run efficiency by transparently interweaving operation and reasoning. Operation contributes to short term efficiency. Reasoning contributes information which can be used to improve future efficiency. The basic task of the planner is to translate job requests into executable operational plans under the objective of long-run efficiency. To achieve the maximum long-run efficiency the planner needs to trade off between short term objectives such as peak performance and long-run objectives such as up-time. The basic task of the reasoning engine is to update the belief state and provide guidance to the planner in terms of information gathering. Both the planner and reasoning engine operate with a common model of the system.

2.3 Current work: Two concreate instantiations of *Pervasive Reasoning*

In recent work I instantiated the framework of pervasive reasoning in two ways:

- Pervasive Diagnosis: Integration of Active Diagnosis into Production Plans (Kuhn *et al.*, 2008), where I'm interested in stimulating the planner to gather information to better learn failure probabilities of actions.
- Pervasive Model Adaptation: Integration of Planning and Information Gathering in Dynamic Production Systems (Liu *et al.*, 2009 to apear), where I'm interested in stimulating the planner to gather information to better adapt action parameters in an evolving system environment.

In both cases I implemented a heuristic search approach that guides the planner to use plan flexibility to increase information gain or to trade off productivity with information gain. In the prototype, the heuristic search is embedded in an existing model-based planner (Do *et al.*, 2008) and combined with a diagnosis/model-adaptation engine which provides continuous belief updates to the planner. Experimental results show that pervasive reasoning reduces the loss rate and improves the long-run productivity over passive and explicit reasoning.

2.4 Proposed Work: Extend *Pervasive Reasoning* to learn domain structure and policies

My current work enables systems to actively manipulate operation to gather more information to improve learning of probabilities as well as continuous parameters. I see another application of automated health and status information gathering in prognosis, which I call pervasive prognosis. Similar to pervasive diagnosis, pervasive prognosis actively manipulates production plans to maximize information gain. Due to those manipulations the system can be exercised such that the feedback allows a more accurate estimation of the remaining life-time.

3 CONCLUSIONS

The idea of Pervasive Reasoning opens up new opportunities to efficiently exploit information for the optimization of model-based systems. Reasoning about intermittent events which would have required expensive production stoppages (for explicit reasoning) can now be addressed online during production. While pervasive reasoning has interesting theoretical advantages, empirical results show that a clever combination of heuristic planning and reasoning can be used to create practical real time applications as well.

REFERENCES

(Do *et al.*, 2008) Minh Do, Wheeler Ruml, and Rong Zhou. On-line planning and scheduling: An application to controlling modular printers. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI-08)*, 2008.

(Kuhn *et al.*, 2008) Lukas Kuhn, Bob Price, Johan de Kleer, Minh Do, and Rong Zhou. Pervasive diagnosis: the integration of active diagnosis into production plans. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI-08)*, 2008.

(Liu *et al.*, 2009 to apear) Juan Liu, Lukas Kuhn, and Johan de Kleer. Pervasive model adaptation: Integration of planning and information gathering in dynamic production systems. In *Proceedings of ICAPS-09*, 2009, to apear.