# Towards a POMDP-based Intelligent Assistant for Power Plants

**Alberto Reyes Ballesteros**
Instituto de Investigaciones Eléctricas
Cuernavaca, México 62490

**Matthijs T. J. Spaan**
Institute for Systems and Robotics
Instituto Superior Técnico
Lisbon, Portugal

**Abstract**

This extended abstract presents a decision support system based on decision-theoretic planning techniques. Its goal is to provide power plant operators with useful recommendations to (i) maintain a plant running under safe conditions, and (ii) to deal with process transients when an unexpected event occurs. We use the formalism of partially observable Markov decision processes as the core of an intelligent assistant for combined-cycle power plants, comment some preliminary results, and discuss some ideas about how to deal with uncertain plant states.

## Introduction

In many industrial processes, plant operators are faced with a large amount of problems and information based on which they have to make decisions. To support such decisions, automatic assistants that provide operators with a list of suggested commands can be used (Reyes et al. 2003). Either when the operator executes a command or when a unexpected disturbance occurs, a new list of recommended actions is presented. This recommendation process can be modelled as a sequential decision problem under uncertainty with an optimization criteria such as performance, availability, reliability, or security. Thus, we suggest the use of Partially Observable Markov Decision Processes (POMDP), a well known stochastic method for sequential decisions.

In the context of our application, intelligent assistants (IA) are knowledge-based systems for the decision support that provide suggestions and criticisms during the decision making process (Aamodt and Nygard 1995). In this work, we present a problem where an intelligent assistant should provide operators with on-line guidance in the form of ordered recommendations. The system will allow to deal with abnormal situations, unexpected events, or the occurrence of process transients under uncertainty conditions.

## Problem Domain

In order to illustrate how important the decisions of a human operator are to overcome problems or to optimize operations, we have selected an electric load disturbance in the steam generation system of a combined cycle power plant.

ICAPS 2010 POMDP Practitioners Workshop, May 12, 2010, Toronto, Canada.

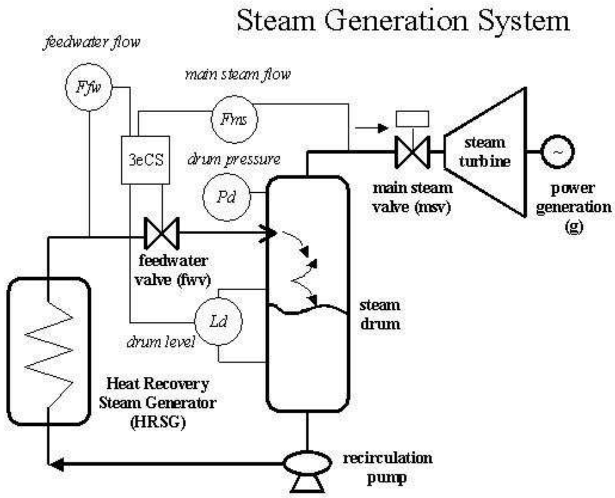


Figure 1: Simplified diagram of the steam generation system. The gas turbine exhaust is not shown. Thin lines denotate instrument connections to pipelines and process equipment (bold lines)

A heat recovery steam generator (*HRSG*) is a process machinery capable of recovering residual energy from the exhaust gases of a gas turbine to heat up water. Eventually, the outlet of the HRSG is a mixture of steam and water (*Ffw*) flowing to the *steam drum* through the feedwater valve (*fwv*). The steam drum is a special vessel whose main function is the efficient separation of mixture that guarantees dry steam flow to the *steam turbine*. This steam flow is regulated by the main steam valve (*msv*). The *recirculation pump* is a device that extracts residual water from the steam drum to keep a water supply in the HRSG. The result of this process is a high-pressure steam flow (*Fms*) in the *steam turbine* that produce electric energy ($g$) in a *power generator*. The complete process control domain is shown in Fig. 1.

An electric load rejection ($d$) is an exogenous event caused by a sudden load disconnexion that could induce abnormal state transitions in the process. Under these circunstances, the current control systems are not efficient and human intervention is required.

A practical solution is the use of an intelligent operator assistant providing recommendations about how to correct the problem. The operator assistant should be able to find an action policy according to the crisis dimension, take into account that sensors are not perfect and that actuators can produce undesired effects. Furthermore, it should consider the performance, availability and reliability of the actual plant installations under these situations.

## Partially Observable Markov Decision Processes

A partially observable Markov decision process (POMDP) (Kaelbling, Littman, and Cassandra 1998) models a sequential decision problem, in which a system evolves over time and is controlled by an agent. At discrete time intervals the agent partially observes the state of the system and chooses an action. In a POMDP there is a set of observations labels $O$ and a set of conditional probabilities $P(o|a, s)$, $o \in O$, $a \in A$, $s \in S$, such that if the system makes a transition to state $s$ with action $a$, it receives the observation label $o$ with probability $P(o|a, s)$. A standard technique for finding an optimal policy for a POMDP is to construct an MDP whose states are the belief state of the original POMDP; that is, each state is a probability distribution over states in the POMDP, with beliefs maintained based on the observation labels using Bayes's rule. A form of value iteration can be performed in this space using the fact that each finite-horizon policy will be convex and piecewise linear.

## Preliminary results and discussion

We have used an MDP-based intelligent assistant (Reyes, Spaan, and Sucar 2009) to run a series of preliminary experiments. In this work, the control strategy followed the idea of looking for plant states with the highest utility that, when analyzed by an expert operator, was near-optimal in most of the state space. A drawback of an MDP model is that is assumes that all the state variables relevant for decision making are observed without noise, while real-world sensors are prone to noise. Furthermore, there might be situations that cannot be detected directly using the available sensors. In this case the state observed by the MDP is no longer Markovian, and hence the value of the computed policies will no longer be accurate.

For instance, during normal operation, the conventional three-element feedwater control system ($3eCS$) commands the feedwater control valve ($fwv$) to regulate the steam drum level ($Ld$). However, when a partial or total electric load rejection is presented this traditional control loop is not longer capable to stabilize the drum level. In this case, the steam-water equilibrium point changes, causing an enthalpy change of both fluids (steam and water). Consequently, the enthalpy change causes an increment in the water level because of a strong water displacement to the steam drum. The control system reacts by closing the feedwater control valve. However, an increment of feedwater is needed instead of a decrement.

Also, the mood or emotional state of an operator is a very important variable which is not observed directly. After all, in a decision support system, it is the operator who finally closes the control loop (e.g., who executes the recommended action).

To tackle at the same time the problem of noisy sensors and limited observability, we propose using POMDPs. Efficient algorithms for approximately solving factored POMDPs are available (Poupart 2005), and no global changes to the original MDP-based operator assistant architecture have to be made.

We can model the load rejection problem detailed above by adding a (binary) state variable that models whether or not a load rejection is occurring. This variable cannot be observed directly, but we maintain a belief whether its true or false. This belief can be updated (using Bayes' rule) given observations of the other state variables, as they can give a clue about its state, for instance because actions do not have the intended effect. In a POMDP formulation such clues can be directly coupled to the state of the load-rejection variable. Adding the variable in a POMDP setting will allow the system to consider the possibility of a load rejection, and to optimize the policy in case it happens, even if it cannot be detected directly. Also, observations regarding the state of the operator can be included in this model.

To conclude, these techniques can also be successfully applied to similar domains. For instance, either in power industry or petroleum industry there are a number of different chemical processes to control optimally that cannot be detected directly. Among them we could mention: oil refinement, water treatment, cooling systems, and oil production. In electric power systems, the electric distribution process is other that could be effectively solved using POMDPs.

**Acknowledgments** This work was partially supported by Fundação para a Ciência e a Tecnologia (ISR/IST plurianual funding) through the PIDDAC Program funds, project PTDC/EEA-ACR/73266/2006, FONCICYT project 95185, as well as by IIE through project 13773.

## References

Aamodt, A., and Nygard, M. 1995. Different roles and mutual dependencies of data, information and knowledge: An AI perspective on their integration. *Data and Knowledge Engineering (North-Holland Elsevier)* 16:191–222.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101:99–134.

Poupart, P. 2005. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. Ph.D. Dissertation, University of Toronto.

Reyes, A.; Sucar, L. E.; Morales, E.; and Ibarguengoytia, P. H. 2003. Abstract MDPs using qualitative change predicates: An application in power generation. In *Planning under Uncertainty in Real-World Problems Workshop. Neural Information Processing Systems (NIPS-03)*.

Reyes, A.; Spaan, M. T. J.; and Sucar, L. E. 2009. An intelligent assistant for power plants based on factored MDPs. In *IEEE Int. Conf. on Intelligent System Applications to Power Systems*.