12th IMEKO TC1 & TC7 Joint Symposium on
Man Science & Measurement
September, 3 – 5, 2008, Annecy, France

OPTIMAL MEASUREMENT SCHEDULING AND DESIGN WITH DYNAMIC PROGRAMMING

Aino Ropponen [1], Risto Ritala[2],

[1] Tampere University of Technology, Tampere, Finland, aino.ropponen@tut.fi
[2] Tampere University of Technology, Tampere, Finland, risto.ritala@tut.fi

Abstract: We shall present approaches to optimize the design of a measurement system and schedule dynamically a versatile measurement resource. The analysis is based on expressing the system management task as a dynamic programming problem in which the system state is partially observable. We shall review the well-known linear-quadratic-Gaussian case, and discuss and give examples solutions of discrete state systems. Furthermore, we discuss on approximate dynamic programming methods to solve such problems in practice.

Keywords: Optimal measurement, Dynamic programming, Decision support.

1. INTRODUCTION

Measurement selection is an important task of system design. As the decisions to operate the system are made based on information available about the system, it is not trivial to integrate design of measurements to overall system design, in particular when measurement subsystem constitutes a considerable amount of overall system costs. One cost-efficient option of measurement subsystem design in many applications is a multipurpose analyzer. However, choosing such a device leads to a complex dynamic optimization problem of scheduling which of the possible measurements is made at each time instant. The scheduling problem was first phrased and solved in a linear-quadratic system already by Meier, Peschon and Dressler [1] in 1967. However, in more complicated cases the problem is computationally extremely heavy, and has only recently attracted interest in robotics [see e.g. 2], and also in quality control at industrial processes [3, 4]. Analyzing the value of information through dynamically optimized system performance, measurements can be designed and scheduled so that the limited measurement resources provide the most valuable information.

Current industrial problems in which a systematic approach on measurement scheduling and design provides a great potential of improvement include design laboratory activities supporting on-line measurements and controlling scanning measurements. E.g. in paper industry the quality of the web produced is measured with a scanning device. At present the scanner is operated regularly across the web, but it can be shown that there exist circumstances at which other scan paths will provide a better uniformity of quality [5].

This paper is organized as follows. Section 2 defines the measurement scheduling problem as a dynamic optimization problem. Section 3 discusses both the exact and approximate solution of this problem for discrete state systems. Furthermore Section 3 shows how linear-quadratic-Gaussian case separates into independent measurement scheduling and control problems in this formalism, which was shown already in [1] in somewhat different formulation. Section 4 outlines, how the measurement design problem is solved with the scheduling problem as a subtask. Section 5 shows three examples: the exact solution of a binary state system, and approximate solutions of a three state problem and a five state problem, the latter inspired by a paper machine quality management problem based on laboratory measurements. Section 6 summarizes the main findings and outlines future research.

2. DYNAMIC PROGRAMMING WITH UNCERTAIN MEASUREMENTS – THE SCHEDULING PROBLEM

The system management task is formulated in terms of system states x_t, the control actions u_t and the measurement choice m_t. The measurement choice is to be understood as any combination of simultaneous measurements that the present measurement system M allows. Given the operational cost, at any time as $h(x_{t+1},u_t,m_{t+1})$ the action optimization with time horizon T and discounting factor $0<\alpha<1$ is:

$$u^*\left(p^{(+)}(x_0)\right)=\underset{\{u_t\}_{t=0}^{T-1},\{m_t\}_{t=1}^{T-1}}{\arg\min} E\left\{\sum_{t=0}^{T}\alpha^t h\left(x_{t+1},u_t,m_{t+1}\right)\right\} \quad (1)$$

where $p^{(+)}(x_0)$ is the current state information expressed as state probabilities at time t=0, and after the measurement at t=0 has been made (posteriority denoted by ‘+’). The expectation value $E\{\}$ is to be calculated with respect to the current state information. Hence u^* is a functional of probability distribution if state is continuous, and a function in a d-1 dimensional space if state consist of d discrete states. In this paper we concentrate on discrete state systems. Hence $p^{(+)}(x_0)$ is a d-dimensional vector with additional constraint that its values are semipositive and add up to one.

We assume the system to be a Markov chain with action dependent transition probabilities, i.e. a Markov decision process. Thus the system dynamics is described with state transition matrices $p_{ij}(u)$ that give the probability of state j at next time step provided that present action u has been taken and the present state is i. The measurements are described with conditional probabilities $q^{(m)}(z_t^{(m)}|x_t)$: the probability that with measurement m a measurement result $z_t^{(m)}$ is obtained if the true state is x_t.

Let us assume firstly that the state information after measurement at time t is $p^{(+)}(x_t)$; secondly, that an action u_t is taken; and thirdly, a measurement m_{t+1} made with a result of z_{t+1}. Then by probability propagation of the Markov chain and the Bayesian interpretation of measurement gives:

$$p^{(+)}(x_{t+1}=j\,|\,u,m,z_{t+1}^{(m)},p^{(+)}(x_t)) = \sum_i \frac{q^{(m)}(z_{t+1}^{(m)}\,|\,x_{t+1}=j)\,p_{ij}(u)\,p^{(+)}(x_t=i)}{P(z_{t+1}^{(m)})} \qquad (2)$$
$$P(z_{t+1}^{(m)}) = \sum_{i,j} q^{(m)}(z_{t+1}^{(m)}\,|\,x_{t+1}=j)\,p_{ij}(u)\,p^{(+)}(x_t=i)$$

The latter is the probability of getting the measurement result $z_{t+1}^{(m)}$ when previous state information is $p^{(+)}(x_t)$. Then the optimization problem of Eq.(1) can be rewritten as an iteration (a dynamic programming problem) with V_T being the optimal T-horizon cost:

$$V_T(p^{(+)}(x_t)) = \min_u \min_m \Big[E_{x_t}\{h(x_{t+1},u,m)\} + \alpha E_{z_{t+1}^{(m)}}\{V_{T-1}(p^{(+)}(x_{t+1}\,|\,u,m,z_{t+1}^{(m)},p^{(+)}(x_t)))\}\Big] \qquad (3)$$

If the measurement system is not of multiuse, the minimization with respect to m_{t+1} is trivial and the problem reduces to the standard form of partially observed Markov decision process (POMDP), see e.g. [2].

In the infinite horizon case the iteration of Eq. (3) turns into a functional equation:

$$V(p^{(+)}(x_t)) = \min_u \min_m \Big[E_{x_t}\{h(x_{t+1},u,m)\} + \alpha E_{z_{t+1}^{(m)}}\{V(p^{(+)}(x_{t+1}\,|\,u,m,z_{t+1}^{(m)},p^{(+)}(x_t)))\}\Big] \qquad (4).$$

The T-horizon measurement choice problem at time $t=0$ given state information prior $p^{(-)}(x_0)$ to measurement is then formulated as:

$$m_T^*(p^{(-)}(x_0),u_{-1}) = \arg\min_{m_0}\Big[E_{x_0^{(-)}}\{h(x_0,u_{-1},m_0)\} + E_{z_0^{(m)}}\{V_T(p^{(+)}(x_0\,|\,m,z_0^{(m)},p^{(-)}(x_0))\}\Big]$$
$$p^{(+)}(x_0=j\,|\,m,z_0^{(m)},p^{(-)}(x_0)) = \frac{q^{(m)}(z_{0|}^{(m)}\,|\,x_0=i)\,p^{(-)}(x_0=j)}{\sum_i q^{(m)}(z_{0|}^{(m)}\,|\,x_0=i)\,p^{(-)}(x_0=j)}$$
$$P(z_0^{(m)}) = \sum_i q^{(m)}(z_{0|}^{(m)}\,|\,x_0=i)\,p^{(-)}(x_0=j) \qquad (5).$$

If the measurement cost is additive to the cost of state and action, the optimal measurement choice will not depend on the previous action u_{-1} other than through $p^{(-)}(x_0)$.

In case of infinite horizon, V_T is replaced by V. This completes the definition of the measurement scheduling problem.

3. METHODS TO SOLVE THE MEASUREMENT SCHEDULING PROBLEM

3.1 Exact solution of discrete state case

Sondik proved in 1971 [6,7] that the solution to the problem in the Eq.(3) can be reformulated as

$$V_t(\bar{p}) = \min_{\alpha\in\Gamma_t} \alpha^T p \qquad (6)$$

where $[\bar{p}]_i = p(x_t=i)$, with $\sum_i [\bar{p}]_i = 1$ and α is a $|x|$-dimensional vector and Γ_t is a set of α- models. Therefore the exact solution of Eq. (3) is always piecewise linear and concave in probability $\bar{p}$. Thus each of the α model vectors define the optimal actions, in our case the control action and the choice of measurement, for a certain region of probability $\bar{p}$. The solution to the problem for a given time horizon t can then be presented as the collection of α – models, Γ_t.

The recursion of α –model collection Γ_t from the collection Γ_{t-1} is presented in [2]. The method is rather straightforward. The problem is that as the time horizon increases, the number of α -models increase rapidly, more than exponentially. The size of the problem can be decreased by pruning α – models which do not contribute to optimality for any probability $\bar{p}$. The method to find these pruned vectors is, however, time-consuming and hard.

This exact solution can be directly only used for small problems with few states and/or short time optimization horizon.

3.2 Point-based approximate solution for discrete case

The exact solution optimizes the solution over all probabilities. The difficulty is the increasing number of α – models as optimization time horizon increases. The idea of point-based solution is to solve the problem in a fixed set of probability points. This is based on fact that at any probability point only one α –model is active in the minimization of eq. (6). Obviously, the selection of probability points is critical for the point-based method, but if the points are selected properly, the solution is close to exact optimal solution even for high number of discrete states and long time horizon.

Several methods to select points have been presented in [2]. In our case studies points are selected randomly or set in a regular grid.

When the set of probability points is selected, the optimal α -models are calculated for each probability point. The set of α –models can be pruned taking away the identical models. Thus, the number of the α -models is at most the number of probability points. As a result the recursion becomes simpler and faster. The method is presented more specific e.g. in [2].

3.3 Linear-quadratic-Gaussian case

The transition dynamics and measurement description of linear Gaussian case with continuous valued state vector can be written as:

$$\begin{aligned} f^{(dyn)}(x_{t+1} \mid x_t; u) &= N_d\left(A_t x_t + B_t u_t; \Sigma_t^{(dyn)}\right) \\ f^{(meas,m)}(z_t^{(m)} \mid x_t) &= N_{d'}\left(C_t x_t; \Sigma_t^{(meas,m)}\right) \end{aligned} \tag{7}$$

where $N_d(\mu,\Sigma)$ is a d-dimensional Gaussian distribution with mean μ and covariance matrix Σ.

Let us further assume that state, action and measurement costs are additive and that state and action costs are quadratic:

$$\begin{aligned} &h(x_{t+1}, u_t, m_{t+1}) = (x_{t+1} - \hat{x}_{t+1})^T H_X (x_{t+1} - \hat{x}_{t+1}) + \\ &(u_t - \hat{u}_t)^T H_U (u_t - \hat{u}_t) + H_M(m_{t+1}) \end{aligned} \tag{8}$$

If the initial state information is Gaussian, $p^{(+)}(x_t) = N_d(\mu,\Sigma)$, then by induction with Eq. (3) one can show that

$$V_{T-1}(\mu,\Sigma) = G_{T-1}(\Sigma) + (\mu - \hat{\mu}_{T-1})^T H_{T-1}^{(V)} (\mu - \hat{\mu}_{T-1}) \tag{9}$$

Furthermore, when applying this in Eq.(3), one finds that the cost inside the minimization in Eq. (3) is a sum of two terms: one depending on present state covariance matrix and measurement choice, and another one depending on present state expectation value and control action. Thus the optimal choice of measurement depends only on which measurements have been made in past but not on which measurement values were obtained. As a result the choice of measurement schedule for the linear-quadratic-Gaussian case is a policy: as the system runs and measurement values are obtained, no additional information assisting in the choice of future measurements is obtained. Correspondingly, the optimal action depends only on the present state estimate, not through which measurement choices the estimate has been obtained.

The separability of measurement scheduling and control was already proven in [1]. However, we included this short discussion from within the formalism of Section 2, because of the practical importance for the measurement community.

4. MEASUREMENT SYSTEM DESIGN PROBLEM

The optimal performance achievable depends on the measurement system M. If we assume that the cost $h(x_{t+1},u_t,m_{t+1})$ is scaled to real monetary costs, the cost per time unit, when system is operated optimally with a measurement system M, is given as

$$W(M) = \frac{1-\alpha}{1-\alpha^{T+1}} \int_{domain(p^{(+)}(x))} f(p^{(+)}(x)) V_T(p^{(+)}(x)) d(p^{(+)}(x)) \tag{10}$$

Function $f(p^{(+)}(x_0))$ allocates a probability density to facing a control problem with initial information $p^{(+)}(x_0)$; i.e. it is a probability density on state probabilities. The prefactor in Eq. (10) normalizes the sum of discount weighting factors in Eq. (1) to one so that comparing formulations with different time horizon can be directly compared.

The probability $f(p^{(+)}(x_0))$ may be thought of arising from a component of regular behavior of the system and/or of abnormal behavior. Regular behavior means that the system behaves according to the system model $p_{ij}(u)$. This component of $f(p^{(+)}(x_0))$ can simply be obtained simulating the system with optimal control actions and measurement choices according to solution of Eq. (1), and then estimating the probability density with observations of $p^{(+)}(x_0)$ during the simulation. Abnormal component arises due to that the system is subject to unforeseeable disturbances: the system model $p_{ij}(u)$ is not valid at such instants. It is the designers' task to specify the abnormal scenarios and their probability of occurrence. The corresponding component in $W(M)$, Eq.(10), describes the system performance when recovering from abnormality.

In the design the operational performance is weighed against investment cost $C(M)$. If the measurement investment is to be uniformly depreciated in time T_d (in units of time steps in operation), the design problem is

$$\min_M \left[T_d \cdot W(M) + C(M)\right] \tag{11}$$

Similarly, a discounted depreciation problem may be formulated as a Net Present Value problem, see e.g. [8].

5. CASE STUDIES

This section present three simple case studies of joint dynamic optimization of control and measurement action. First a two-state system is briefly analyzed through exact solution, and then the exact and point-based approximate solutions of a three state system are compared. Finally a short simulation study of a five-state system corresponding to a simple quality management case is discussed

5.1 Two-state system

The simplest case to illustrate the method is a two-state system with two action alternatives and two measurement alternatives. Two-state control problem was addressed already in [9].

In our example the states are “good” and “poor”. Being in the “poor” state incurs an additional cost of 0.9 units. The actions are “run as usual”, and “make a correction”. Under the first action the transition probability from “good” to “poor” is 0.3 and that of the opposite transition 0.2. Under the corrective action the “good” state remains with probability 0.9 and the “poor” state turns into “good” with probability 0.8. The additional cost of corrective action is 0.5 units. Furthermore, the state can be measured at an additional cost of 0.04 units, and the probability of measurement giving the erroneous state is 0.05. The discount factor is 0.95.

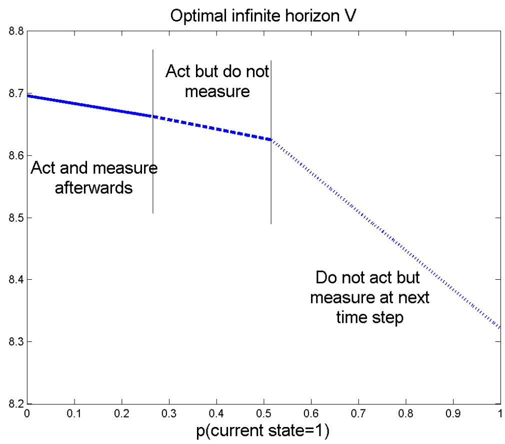

Figure 1. Value function of a two-state system.

Figure 1 presents the value function V and decisions of the infinite horizon case, Eq. (4) as the function of the probability of present state being "good". The infinite horizon solution was obtained by iteration of finite horizon case, Eq. (3) with the exact piecewise linear methods of section 3.1, Eq. (6). If the state is "poor" with high probability, corrective action is made and verified with followup measurement. If there is high uncertainty about the state, corrective action is made, but as a result the probability of "good" state is so high that no followup measurement is needed during the next step but only on further steps. When there is high certainty about state being "good", the state is monitored with measurement, but no corrective actions are needed. However, simpler policies throughout low and medium probability of state being "good" of either both making the corrective action and measurement or only making corrective action are only slightly suboptimal.

In finite horizon case, the horizon affects strongly the policy. With horizons 1 and 2, no measurements are ever made. With horizon 3, measurements are always made. From horizon 4 on, the decision policy is as in infinite horizon case with decision limits converged down to 0.001 accuracy in *p(state=1)* at horizon 7. The average level of V, not relevant to decisions, converges only after horizon 40, due to discount factor being close to 1.

The exact solution in this case is simple as the number of piecewise linear components considered in V remains small, at most 24 if models are pruned at each step. In hindsight, the identical result would have been obtained with the approximate method of Section 3.2 minimally with three points (e.g. 0, 0.4 and 1). If points are chosen at random, or uniformly, the probability of obtaining the exact solution for the binary system, with the approximate method is high with only 10...20 points.

5.2 Three state system

Our second example is a system with three states: "good", "acceptable", "poor". The system is controlled by three actions. The first action, the cheapest one, decreases slightly the probability of the "good" state. The second control option turns the system into a better state and the third one turns the system surely into "good".

Similar, but not the same, problem with three states was presented by Smallwood and Sondik in 1973 [7].

The problem is solved using point-based solution, in which a grid of points is selected. The results for time horizon $T = 2$, $T = 3$ and $T = 5$ are presented in the Figures 2, 3 and 4. Grid points are shown with asterisk-marks in the Fig. 2. Decision borders are shown by lines.

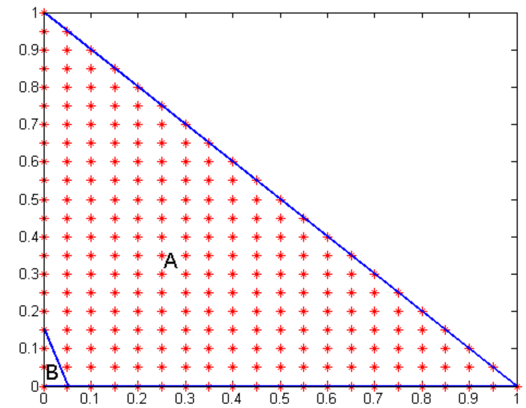

Figure 2. Decision planes when time horizon $t = 2$. Grid points are shown with asterisk-marks and decision borders by lines. In the area A the optimal decision is to make the control action 1 and not to measure. In the area B the optimal decision is to control 2 and not to measure.

At the end of the horizontal axis the probability vector is [1 0 0], the system is certainly at "good" state. At the end of the vertical axis probability vector is [0 1 0], the system is certainly at "acceptable" state, and in origin [0 0 1], the system is certainly at "poor" state.

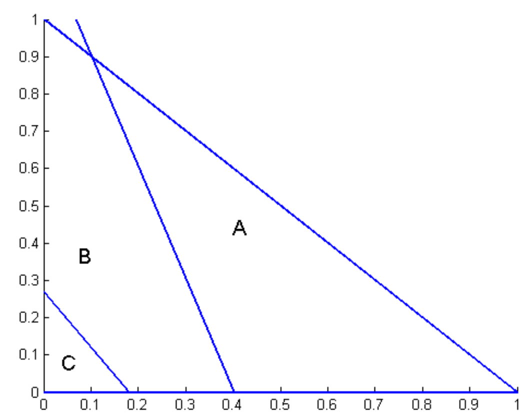

Figure 3. Decision planes when time horizon $T = 3$. In the area A the optimal decision is to make the control action 1 and not to measure. In the area B the optimal decision is to control 2 and not to measure. In the area C the optimal decision is to control 3 and not to measure.

For time horizon T = 2 only two decisions exists: to control 1 or 2 and never to measure. For time horizon t = 3 number of decisions is increased to three, but measuring still never measuring. That is logical as measuring costs and for short decision horizons it is reasonable to only act.

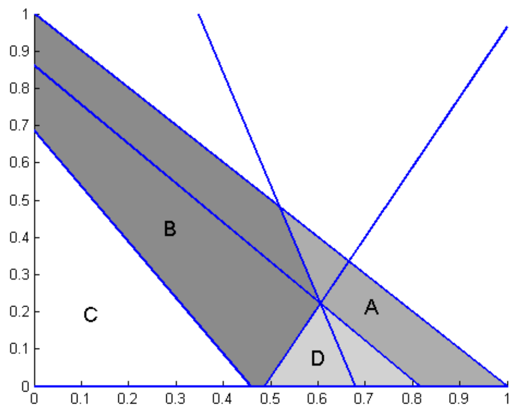

Figure 4. Decision planes when time horizon T = 5. In the area A the optimal decision is to make the control action 1 and not to measure. In the area B the optimal decision is to control 2 and not to measure. In the area C the optimal decision is to control 3 and not to measure. In the area D the optimal decision is to control 1 and to measure.

For time horizon T = 5 the number of decision planes is increased to four, with "control 1 and measure" as a one option as shown in the Fig. 4. Precisely the same result is obtained with the exact solution method.

As the time horizon grows the area of control action 3 increases and the area of control action 1 decreases. That is logical as in the long term it becomes more important to ensure the better result and with control action 3 the system is surely turned into "good" state. Also the importance of measuring grows with the longer time horizon.

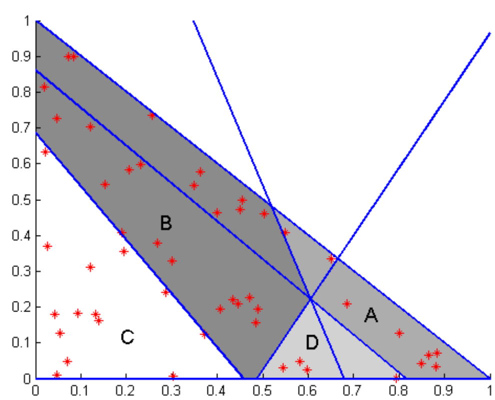

Figure 5. Decision planes when time horizon t = 5 and the points are selected randomly (50 points).

The same result, i.e. the same decision planes, is achieved also using 50 uniformly distributed random points as shown in the Fig. 5. As only four points - one from each plane - is needed to obtain the optimal result, the number of random points could be less and still that probability to obtain the optimal solution would be high enough.

As it can be seen from the figures, horizon affects strongly the decision borders and the optimal decisions. With higher time horizons (T>6), the number of decision borders further increases and the figures are no longer illustrative as can be seen in the Fig. 6 where the horizon T=8. There are 10 decision planes, but as the border lines cross, the planes are not easy to see. Despite the illustration problem the calculation is still straight forward and quick to compute.

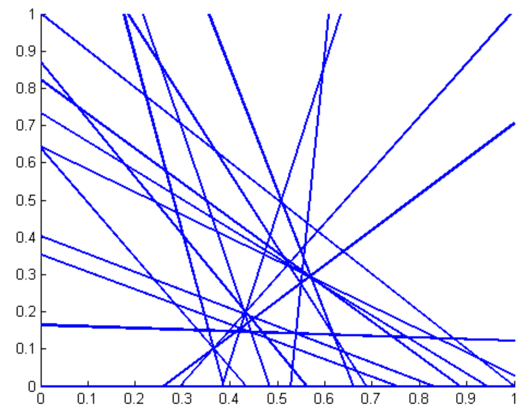

Figure 6. Decision borders for time horizon T = 8 using point-based approximate method with points selected from the grid.

The solution obtained by the point-based method with T = 8 differs slightly of the solution obtained using exact method where 12 decision planes exist (Fig. 7). However, as the missing planes are rather small and as the calculation time using exact method is 400-fold it is justified to use point-based approximate method instead.

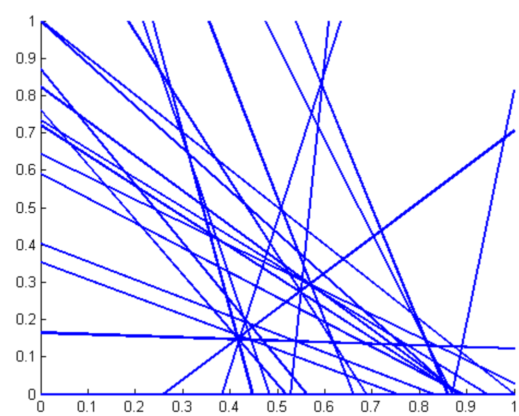

Figure 7. Decision borders for time horizon T = 8 using exact solution method.

5.3 Paper machine quality control case

As a five-state model we present an example which is inspired by a quality management problem in papermaking.

Table 1. Example of simulated result from the paper quality case. The control and measurement actions are optimized with time horizon of *T*=4. The first line shows the measurement decision with measured value or '-' denoting not a measurement. The second line gives the optimal control action. The third line gives the true process state that is not observable to optimizer. Note the different scale of measurements and states.

	t=0	t=1	t=2	t=3	t=4	t=5	t=6	t=7	t=8	t=9	t=10	t=11	t=12	t=13	t=14
Measured value	-	-	14	16	-	16	-	13	11	-	-	13	11	-	-
Optimal control	4	3	3	1	3	1	3	3	4	1	3	3	4	1	3
True value	2	4	5	5	5	5	4	3	4	5	4	4	4	5	5

The target of the case is to control quality variable tear strength optimally. Five states have been defined for tear strength, 1 referring to “poor” state and 5 to “good” state. The costs of quality at states 1 to 5 are 10, 5, 2, 1 and 0 units.

Tear strength is controlled by fiber furnish fraction ratio which is discretized into four values. The costs of control actions from 1 to 4 are 1, 2, 3 and 4 units. Measurement options are again to measure or not to measure, with costs 0.3 and 0 units. Thus, 8 action alternatives exist.

The case is simulated by running two simulators in parallel. The first simulator optimizes the control and measurement actions and acts according to them. The second simulator is the so called “true state simulator” which calculates the evolution of the target system. The true state is not known by the first simulator, as the state is known only through uncertain measurement of the second simulator.

The information about the state of the quality, the probability vector, is updated after each control and measurement action. The control and measurement actions are optimized based on the information about the current quality state.

As the point space in a five-state system is wide, it is difficult to assess the correct number of points and optimal solution is hard to find. Using 100 points, the number of solution planes on the first run is around 30. Using 1000 points the number of planes is around 100 and using 10000 points around 230. However, even though the solution differs with different number of points, the quality is manageable with fewer points also. Good suboptimal results can be achieved by using only 100 points.

An example of results is shown in the Table 1. The problem is solved using point-based solution with 500 points. The corner points are selected beforehand and other points randomly from a uniform distribution. The optimization time horizon in this simulation is 4 time steps. At the beginning the quality is at a state 2, but recovers to good quality quite quickly and stays in good quality.

6. DISCUSSION AND FUTURE WORK

We have presented the measurement scheduling and design problems of a stochastic dynamic system in terms of dynamic programming. We have provided examples in which both exact and point-based approximate methods to solve the dynamic optimization problem have been applied. We have presented the solution of the problem in the simplest of cases, a binary state system, and in the three-state system. The presented five-state system follows an industrial quality management case [3,4].

An obvious problem of the point-based approximate solution is the selection of points. It cannot be guaranteed that the optimal solution is found using the approximate solution. On the other hand using the exact solution the calculation time grows extremely rapidly as time horizon increases.

The dynamic programming approach relies heavily on the system model $p_{ij}(u)$. This raises – as in all model-based control or optimization problems – the issue of robustness against model inaccuracy. Such robustness problems can be addressed with the well-known method of Q-learning in infinite horizon dynamic programming problems [10, 11]. However, applying Q-learning in the case of uncertain measurements has not been presented in literature.

At present, we are working on a practical application at papermaking: the overall quality control, in particular managing paper strength and brightness. This control is based on laboratory measurements that are rather uncertain as only few paper sheets are taken to represent the machine reel of 40 metric tons of paper. Furthermore, the effects of the main control actions – furnish component ratio and dosage of bleaching chemicals – are known somewhat vaguely due to nonlinearities, complex interactions and long dead times. Therefore, the present description that discretizes the quality parameters, measurement results and control actions and presents the system dynamics through conditional probabilities appears appropriate.

The decisions about strength/brightness control actions are made by several operators and currently there is little if any communication to make the actions coherent. The optimization approach outlined in this paper may serve in addition to automated quality management tool as a operations’ decision support tool for harmonizing the operator actions.

In Section 4 we presented a systematic approach for designing measurement systems. Although there are obvious difficulties in applying it in practice – e.g. defining the control scenarios – the approach puts the process system dynamics design and measurement/control design on an equal footing. Currently we are studying how this approach can be applied to concurrent design of material and information flows of a paper production system.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the financial support provided by the Academy of Finland, through research contract on project "Systematic design of measurement information systems – SyDeMIS".

REFERENCES

[1] L. Meier, J. Peschon, R.M. Dressler, "Optimal Control of Measurement Subsystem", IEEE Transactions on Automatic Control, vol. AC-12, no. 5, 528-536, October 1967.
[2] J. Pineau, G. Gordon, S. Thrun, "Anytime point-based approximations for large POMDPs", J. Artificial Intelligence Res. 27, 335-380, 2006.
[3] A. Ropponen, R. Ritala, "Production-line wide dynamic Bayesian network model for quality management in papermaking", in B. Braunschweig, X. Joulia, editors, Proceedings of 18th European Symposium on Computer Aided Process Engineering, pp 979-984, Lyon, France, June 2008.
[4] A. Ropponen, R. Ritala, "Towards coherent quality management with Bayesian network quality model and stochastic dynamic optimization", accepted to Control Systems 2008, Vancouver, Canada, June 2008.
[5] J. Ylisaari, K. Konkarikoski, R. Ritala, T. Kokko, M. Mäntylä, "Scanner path as manipulatable variable for optimal control and diagnostics of CD variations", accepted to Control Systems 2008, Vancouver, Canada, June 2008.
[6] E. J. Sondik, "The Optimal Control of Partially Observable Markov Processes", Ph.D. thesis, Stanford University, 1971.
[7] R. E. Smallwood, E. J. Sondik, "The optimal control of partially observed Markov processes over a finite horizon", Operations Research, 21, 1071-1088, 1973.
[8] L. T. Biegler, I. E. Grossmann, A. W. Westerberg, "Systematic methods for chemical process design", Prentice-Hall, 1997.
[9] K. J. Åström, "Optimal control of Markov processes with incomplete state information", Journal of Mathematical Analysis and Applications, vol. 10, issue 1, 174-205, 1965.
[10] D. P. Bertsekas, J. N. Tsitsiklis, "Neuro-dynamic programming", Athena Publishing, 1996.
[11] W. P. Powell, "Approximate Dynamic Programming: Solving the Curses of Dimensionality", Wiley-Interscience, 2007.