Predictive Probabilistic Models for Treatment Planning in Paediatric Cardiology

Niels PEEK
Dept. of Computer Science, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: niels@cs.ruu.nl

ABSTRACT

The planning of clinical treatment actions for children with congenital heart disease requires a subtle trade-off between their immediate and long-term consequences, where most of these consequences cannot be predicted with certainty. It is described how this problem can be cast as a finite-horizon, partially observable Markov decision process. The complexity of the resulting model is reduced by using a graphical representation of state space and transition probabilities; it is shown that such a representation yields a significant decrease in the number of model parameters that have to be assessed.

1. INTRODUCTION

A significant problem for the paediatric cardiologist in the management of patients with a cardiac anomaly is to decide if and when a patient has to be submitted to surgical treatment. In the management of these patients, there is always a trade-off between the benefits gained by waiting before surgical intervention in the hope that the patient's condition will improve, and the risks caused by the poor natural history of these disorders [1]. The number of factors involved in this decision-making process is large and their interplay is subtle. Therefore, it is extremely difficult for the clinician to determine which timing of medical and surgical treatment will be optimal for a given patient. In general terms, this problem may be characterised as prognostic assessment and planning under uncertainty with time constraints.

Partially observable Markov decision processes (POMDPs) [2], [3] are models for sequential decision making under uncertainty, taking into account both the outcomes of current decisions and future decision-making opportunities. The general form of these models allows for the expression of many different decision-making scenarios, including reasoning with incomplete information, planning of both test and treatment actions, and predicting future states. Markov decision processes are receiving increasing attention from the AI community, and have recently been proposed as a suitable framework for decision-theoretic planning [4]. However, the generality of the framework precludes straightforward application in practice, due to the representational and computational complexity of POMDP problems of considerable size.

In this paper, we investigate the applicability of POMDPs to the problem of time-critical treatment planning and prognosis in the domain of paediatric cardiology. As a case study we have selected the relatively frequent occurring cardiac anomaly of ventricular septal defect. In this problem domain, truthful modelling of clinical treatment practice requires some adaptations to the general form of the POMDP model. Furthermore, we consider the graphical representation of POMDP models using *temporal belief networks* [5], [6]. These networks have previously been suggested as an adequate representation method for biomedical processes over time [7], [8], and, more recently, as a suitable means to reduce the representational complexity of POMDP models [4]. The case study reported in this paper provides evidence to support this claim.

This paper is organised as follows. In Section 2, we discuss the problem of clinical treatment planning for children with a ventricular septal defect. Section 3 reviews Markov decision process models and associated solution techniques. In Section 4, we discuss the representation of Markov decision processes in probabilistic networks. The resulting representation method is analysed in Section 5 in terms of representational complexity. The paper is completed with conclusions and directions for further research in Section 6.

2. EXAMPLE: TREATMENT PLANNING FOR VSD

Ventricular septal defect (VSD) is a relatively well-understood cardiac anomaly with many clinical features that are typical for congenital heart disease in general. VSD is a defect in the ventricular septum, the fibromuscular wall that separates the left and the right ventricle. An immediate consequence of this defect is blood flow ("shunt") from the left to the right ventricle due to ventricular pressure differences. The shunt size, *i.e.*, the amount of blood flowing through the defect, depends primarily on the size of the defect and the ratio of pulmonary and systemic vascular resistances. Left-to-right shunting causes oxygenous blood to be pumped into the lungs again. As a result, the pulmonary pressure will rise, and systemic cardiac output will decrease. With large defects, the high pulmonary pressure may lead to heart failure. Heart failure accounts for most of the typical symptoms associated with VSDs, such as shortness of breath, feeding problems, oedema, and growth arrearage.

The clinical course is favourable with small defects

To be presented at the Invited Session on *Intelligent Prognostic Methods in Medical Diagnosis and Treatment Planning* during *Computational Engineering in Systems Applications* (CESA'98), Nabeul-Hammamet, Tunisia, April 1998.

throughout infancy and childhood [9]. About 75 to 80% of the defects close spontaneously due to tissue growth, with the majority closing in the first two years of life. Patients with moderate-sized defects may develop large left-to-right shunts and associated complications in infancy, but the majority of this group can be managed medically without surgical intervention. Patients with large defects are more difficult to manage, because of the risks of mortality in the first year of life due to heart failure and associated pulmonary infections. Furthermore, elevated pulmonary vascular resistance may develop over time as a response to continuous pulmonary overflow and hypertension [10]; this is termed *Eisenmenger's complex*. As a result, the shunt size will decrease and the symptoms of heart failure will vanish. However, Eisenmenger's complex eventually leads to cyanosis due to shunt reversal, and severe, irreversible damage to the pulmonary arteries (*arteriopathy*). Early surgical intervention is therefore strongly recommended for these patients. The majority of patients with repair of uncomplicated VSD in infancy or early childhood have an excellent result with no clinical signs or symptoms and apparently normal life-expectancy.

Variable	*Interpretation*	*Domain*
art	pulmonary arteriopathy	*absent*, *mild*, *moderate*, *severe*
cyan	central cyanosis	*false*, *true*
fail	heart failure	*absent*, *mild*, *moderate*, *severe*
res	relative pulmonary vascular resistance	*normal*, *increased*, *high*, *very_high*
risk	mortality risk	*false*, *true*
shunt	shunt size	*none*, *small*, *large*, *reversed*
size	defect (VSD) size	*null*, *small*, *moderate*, *large*
symp	heart failure symptoms	*absent*, *mild*, *moderate*, *severe*

TABLE I
STOCHASTIC STATE VARIABLES FOR THE VSD DOMAIN.

In this paper, we study the clinical problem of treatment planning for VSD patients. We will use a set of eight, discrete, stochastic variables to characterise the disease process; these variables, and their possible value sets, are listed in Table I. For some of these variables (e.g., shunt size, pulmonary vascular resistance, cyanosis, and heart failure symptoms), at each point in time the actual state depends solely on the actual states of other variables. For other variables (e.g., defect size and pulmonary arteriopathy), the state is determined by their own and other variables' states in the past. Furthermore, the clinician has choice among a number of actions to change and/or observe the states of some variables. For instance, medical treatment may be used to control heart failure, surgery may be used to close the defect, and ultrasound imaging, cardiac catheterisation, and pulmonary biopsy provide additional information concerning the defect size, the shunt size and the presence of pulmonary arteriopathy, respectively. The treatment actions, their effects to state variables, and their associated observations are summarised in Table II.

Action	*Interpretation*	*Effects*	*Observations*
biop	pulmonary biopsy	–	art
cath	cardiac catheterisation	risk	shunt
echo	ultrasound imaging	–	size
med	medical treatment	fail	–
surg	perform surgery	size, risk	size

TABLE II
AVAILABLE TREATMENT ACTIONS FOR THE VSD DOMAIN.

The problem for the clinician is to decide if and when to conduct one or more of these actions for a VSD patient. Usually, the patient's condition is monitored without surgical intervention during the first year of life. During this period, ultrasound images are made repeatedly, and medical treatment is given when necessary. After the first year of life, a decision on surgical intervention is made. In cases of doubt concerning the clinical state of the patient, cardiac catheterisation or pulmonary biopsy may be performed prior to that decision. Therapy is considered completed after closure of the defect, either spontaneously or by surgical intervention.

3. MARKOV DECISION PROCESSES

Markov decision processes (MDPs) [11], [12] are models for sequential decision making under uncertainty, which take into account both immediate and long-term consequences of decisions. Basically, the theory assumes that a person, called the *decision maker*, is charged with the responsibility of choosing a sequence of actions in order to influence a stochastic process. The immediate result of each action choice is that the system under consideration evolves to a new state according to a probability distribution determined by the action choice. Furthermore, the decision maker receives an immediate reward reflecting the desirability of the new system state compared to other possible states. The goal is to optimise some function of the overall reward sequence that expresses the decision maker's intertemporal trade-offs. *Partially observable Markov decision processes* (POMDPs) [2] are a generalisation of Markov decision processes which permit uncertainty regarding the system state and allow for state information acquisition depending on action choice. Consequently, in POMDP problems the trade-off between actions does not only concern their immediate and long-term effects, but also their information-gathering properties. The generalisation to POMDPs is significant in problem settings where state uncertainty is a central issue that cannot be discarded; we note that this holds for the VSD domain as discussed in the previous section.

POMDP model

Formally, a POMDP model is a tuple (T, X, A, P, ω, r), where

- T is a set of *decision moments*,
- X is a set of *stochastic state variables*, jointly defining the set of *system states*,
- A is a set of available *actions*,
- P is a set of *transition probability functions*,
- ω is an *observation function*, and
- r is a *reward function*.

The basic form of the model is quite general and the literature on POMDPs provides a plethora of possible choices for each of the elements. The qualifier "Markov" refers to the fact that the transition probability, observation and reward functions depend on the past only through the current state of the system and the most recent action selected by the decision maker. Below, we discuss the appropriate choices for each of the model elements from the perspective of clinical VSD treatment.

The set T of decision moments denotes the points in time where the decision maker is expected to choose an action. We will confine ourselves to the case where T is finite; one then speaks of a *finite-horizon process*. Standard POMDP theory requires us to fix the set T in advance for a given problem domain. However, in the clinical practice of VSD treatment, decision moments are established by the cardiologist in due course. Therefore, we use the set T as a *grid* of decision moments, where each time point $t \in T$ is a potential decision moment. We take $T = \{0, 1, 2, \ldots, N\}$ ($N \in \mathbb{N}$), where $t \in T$ denotes the age of the patient expressed as number of life months. We adopt the convention that no action is chosen at the last potential decision moment: this moment is included for evaluation of the final system state only. The state space of the POMDP model is characterised by the finite set $X = \{X_1, \ldots, X_m\}$ of discrete, stochastic variables. For the VSD domain, we have that $m = 8$; the variables are listed in Table I. To express a joint assignment of values to variables from a set $X' \subseteq X$, the notion of a *configuration* of X' is introduced, which is denoted by $c_{X'}$. The set of all possible configurations of X' is denoted by $C_{X'}$. Consequently, the state space of the POMDP equals $C_X = C_{\{X_1\}} \times \cdots \times C_{\{X_m\}}$. A state space thus characterised by multiple stochastic variables is sometimes called *structured* [4]. Structured state spaces offer several advantages over flat state spaces (where there is only a single state variable), the most prominent of which is the ability to exploit conditional independencies between the variables at hand; we will elaborate on this subject in Section 4. For the set A of available actions in the VSD domain we refer to Table II. Our approach differs from the standard model presented in the literature in that we allow for multiple, simultaneous actions at each potential decision moment $t \in T$. That is, the clinician is expected to choose a subset $\alpha \subseteq A$ of actions to be performed, where choosing the empty set $\varnothing$ is interpreted as a *skip*-action, i.e. taking no decision at that point in time. Using this action, the clinician can make decisions at irregular time points (by skipping intermediate points) and 'fill up' the remaining time when some satisfactory state is reached.

The effects of actions are described by the set $P = \{p_t^\alpha : C_X \times C_X \to [0,1] \mid t \in T, \alpha \subseteq A\}$ of time- and action-dependent transition probability functions, where $p_t^\alpha(c_X^{\text{post}} \mid c_X^{\text{pre}})$ denotes the probability of arriving at state c_X^{post} after performing action set $\alpha \subseteq A$ in state c_X^{pre} at moment t. Observability of the current system state is modelled by the observation function $\omega : A \to 2^X$, where $\omega(a)$ denotes the set of state variables whose current values can be observed by the decision maker when action a has just been performed. For an action set $\alpha \subseteq A$, we define $\overline{\omega}(\alpha) = \cup_{a \in \alpha} \omega(a)$ to be the observable set. The sets of observable state variables for the VSD domain are listed along with the available actions in Table II. Finally, the immediate rewards of actions are described by the function $r : C_{X_r} \to \mathbb{R}$, where $X_r = \{\mathsf{art}, \mathsf{cyan}, \mathsf{risk}\}$; the states of these variables are generally taken to measure success of treatment in the VSD domain, whereas the other variables are considered irrelevant for that matter. The numerical value $r(c_{X_r})$ denotes the desirability of configuration c_{X_r} compared to other possible configurations of X_r. We note that in more general models also the most recent action choice (e.g. associated financial costs) may be taken into consideration by the reward function.

Problem specification and solution

A POMDP problem specification consists of a POMDP model, an initial probability distribution Pr on the state space, and a *utility function* $u : \mathcal{R} \to \mathbb{R}$ reflecting the decision maker's intertemporal trade-offs, where $\mathcal{R}$ is the set of all possible reward sequences for the given model. For the VSD domain, we will use the *total reward criterion*, where $u(r_1, \ldots, r_N) = \sum_{i=1}^{N} r_i$. The *sequential decision problem* is to choose, prior to the first decision, a *policy* to maximise the utility function u, where a policy provides the decision maker with a prescription for choosing an action in any possible state occupied by the system.

Generally speaking, the complexity of finding utility-maximising policies depends on the size of the state space, the number of available actions, and the horizon length. For finite-state (fully observable) MDP problems, efficient solution methods exist, based on the principle of dynamic programming [12]. Unfortunately, this does not hold for problems involving partial observability. Solving a POMDP problem directly necessitates keeping track of entire process histories, of which the sizes grow exponentially in the size of the state space and action sets. A more promising approach is based on transforming the (finite-state) POMDP model to an equivalent MDP model in which system states are probability distributions on the state space of the original POMDP model [2]. When the state space in the original POMDP model contains n states, then the state space in the transformed MDP model is the $(n-1)$-dimensional unit simplex; transition probabilities are derived through Bayes' rule. Dynamic programming techniques can be applied to solve problems using the transformed MDP model, but because of the continuity of its state space, algorithms are complicated and limited. Solving POMDPs with a short, finite planning horizon is nevertheless feasible [3].

4. GRAPHICAL REPRESENTATIONS OF POMDPS

It was noted above that structured POMDP state spaces offer a number of advantages over flat state spaces. One of these advantages is that we can factorise the transition probability distributions describing the effects of actions. Factorisation of a joint probability distribution is based on conditional independence relations induced by the distribution, and allows for a reduction in the required number of model parameters, and for more efficient probabilistic inferences. *Bayesian belief networks* [13] provide for a

concise, graphical representation of factorised joint probability distributions. More recently, the belief network framework was extended to cope with dynamic stochastic systems [5], [6], where the joint probability distribution on the variables in the network evolves over time. These networks are generally referred to as *temporal* belief networks. Temporal belief networks have been suggested as a suitable way to express the dynamics of POMDP models [4], equally facilitating the solution methods by exploiting independencies between state variables, and by making explicit persistence of states and ramification of action effects [5].

Now, let $X(t)$ denote the set of state variables at time point t. A *two-stage temporal belief network* (2TBN) [4] is a belief network with two sets of variables $X(t)$ and $X(t+1)$, where each arc is drawn either a variable from $X(t)$ to a variable from $X(t+1)$, between two variables both from either $X(t)$ or $X(t+1)$. The former type of arc is called a *temporal* dependency; the latter is called an *atemporal*, or *synchronic*, dependency. 2TBNs allow for a compact representation of probabilistic state transitions, where temporal dependencies model the direct effects of actions, and synchronic dependencies model intra-stage correlations. These intra-stage correlations are independent of action choice, so we can use a single, atemporal belief network for their specification.

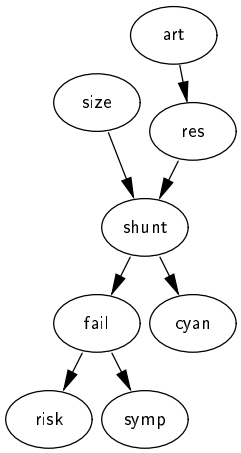

Fig. 1. Synchronic belief network for the VSD domain.

An atemporal belief network for the VSD domain has recently been constructed with aid of a domain expert [14]; a simplified version of the network is shown is Fig. 1. For $X = \{X_1, \ldots, X_m\}$, let , $_s = \{\gamma_s^i : C_{\rho_s(X_i)} \times C_{X_i} \rightarrow [0,1] \mid i = 1, \ldots, m\}$ be the set of synchronic probability assessment functions associated with the belief network, where $\rho_s(X_i)$ denotes the parent set of variable X_i in the synchronic dependency graph; $\gamma_s^i(c_{X_i} \mid c_{\rho_s(X_i)})$ denotes the probability of X_i having configuration c_{X_i} given the configuration $c_{\rho_s(X_i)}$ of its parents. We assume these probabilities to be stationary (i.e. independent of time). The atemporal belief network is also taken to represent the initial probability distribution on the state space in the POMDP problem specification.

We distinguish two kinds of temporal dependency: those induced by action choices (*exogenous change*) and those stemming from persistence or action-independent change of state variables (*endogenous change*). Both kinds of dependency can be modelled using 2TBNs, comprising only those state variables that are directly relevant. Fig. 2a shows the 2TBN modelling endogenous change due to Eisenmenger's complex: pulmonary arteriopathy may result from left-to-right shunts. This process is progressive, and therefore we also take into account the former state of the pulmonary arteries. Fig. 2b shows the 2TBN modelling exogenous change due to surgery. The surgery action itself is depicted by a square box, called an *action node*. Successful surgery results in a closed defect (i.e., size = *null*). However, the chances of success equally depend on the size of the existing defect: large defects are more difficult to patch. The dashed arrow indicates that the clinician is able to observe the defect that results. Furthermore, a small risk is associated with the operation. The twisted arc indicates that there is also a path in the synchronic dependency graph from size to risk.

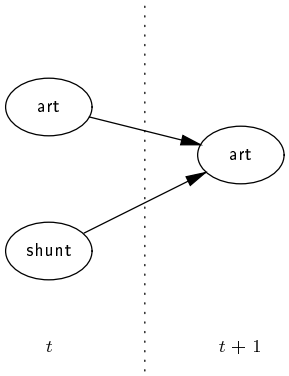

Fig. 2a. 2TBN modelling Eisenmenger's complex.

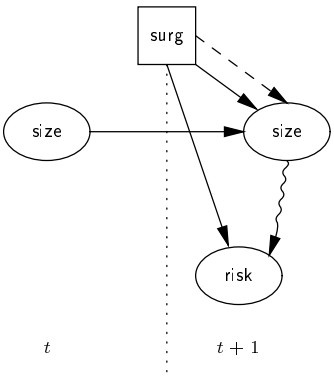

Fig. 2b. 2TBN modelling surgical intervention

If we merge the atemporal belief network and each of the 2TBNs modelling endogenous and exogenous change, a complete graphical model for transitions and observations is obtained. This model is shown in Fig. 3. Formally, let $\rho_t(X_i)$ be the set of temporal predecessors of node $X_i \in X$ in the complete model, that is, there is an arc from $X_j(t)$ to $X_i(t+1)$ for every $X_j \in \rho_t(X_i)$, and let $A_i \subseteq A$ be the set of actions directly affecting X_i, i.e.

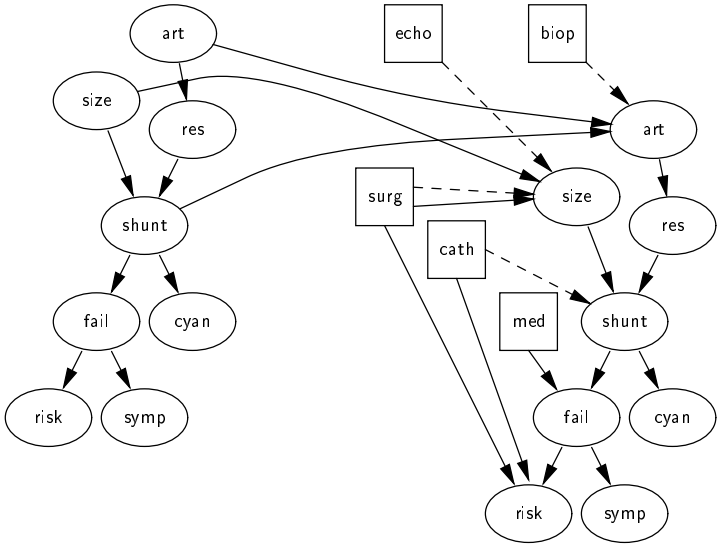

Fig. 3. 2TBN with action nodes for the VSD domain.

there is a solid arc from action node a to X_i for every $a \in A_i$. Furthermore, let , $_t = \{\gamma_t^i : C_{\rho_s(X_i)} \times C_{\rho_t(X_i)} \times 2^{A_i} \times C_{X_i} \to [0,1] \mid i = 1, \ldots, m\}$ be the set of transition probability assessment functions associated with the complete graphical model at time t; $\gamma_t^i(c_{X_i}^{\text{post}} \mid c_{\rho_s(X_i)}^{\text{post}}, c_{\rho_t(X_i)}^{\text{pre}}, \alpha_i)$ denotes the probability of X_i having configuration $c_{X_i}^{\text{post}}$ at time $t+1$ given that its synchronic predecessors then have configuration $c_{\rho_s(X_i)}^{\text{post}}$, its temporal predecessors had configuration $c_{\rho_t(X_i)}^{\text{pre}}$ at the previous moment t, and subsequently action set choice $\alpha_i \subseteq A_i$ was made. We assume that $\gamma_t^i(c_{X_i}^{\text{post}} \mid c_{\rho_s(X_i)}^{\text{post}}, c_{\varnothing}^{\text{pre}}, \varnothing) = \gamma_s^i(c_{X_i}^{\text{post}} \mid c_{\rho_s(X_i)}^{\text{post}})$, i.e., synchronic probability assessment functions are used when there are no temporal predecessors and no action is selected that directly affects variable X_i. In the VSD example, this holds for at least four, and at most six variables, depending on the action set choice. Two variables, size and art are subject to endogenous change connected to the natural course of disease.

The transition probability distributions associated with the POMDP model can now be factorised according to the independency relations portrayed by the graphical part of the complete 2TBN (i.e. with all action nodes added). For each potential decision moment $t \in T$ and each action set $\alpha \subseteq A$, we have that the transition probability function p_t^α can be written as

$$p_t^\alpha(c_X^{\text{post}} \mid c_X^{\text{pre}}) = \prod_{i=1,\ldots,m} \gamma_t^i(c_X^{\text{post}}(\{X_i\}) \mid c_X^{\text{post}}(\rho_s(X_i)), c_X^{\text{pre}}(\rho_t(X_i)), \alpha \cap A_i) \quad (1)$$

for each $c_X^{\text{pre}}, c_X^{\text{post}} \in C_X$, where $c_X^{\text{pre}}(X')$, $c_X^{\text{post}}(X')$ denote the configurations of $X' \subseteq X$ within c_X^{pre} and c_X^{post}, respectively.

5. ANALYSIS

It was shown in the previous section how 2TBNs with action nodes provide for a concise, graphical representation of the dynamics of POMDP models. Concision is arrived at by exploiting both intra-stage and inter-stage independencies between state variables, and explicit representing the limitations in the effects of actions.

Let d be an upper bound on the number of values that each state variable can take, i.e., $|C_{\{X_i\}}| \leq d$ for each $X_i \in X$. So, for $X = \{X_1, X_2, \ldots, X_m\}$, we have that $|C_X| \leq d^m$. Furthermore, let $|A| = k$, so that there are 2^k different action sets. Then, for each $t \in T$ and each $\alpha \subseteq A$, we need at most $d^m \cdot d^m = d^{2m}$ transition probabilities, so the upper bound on total number of parameters needed for the POMDP equals $N \cdot d^{2m} \cdot 2^k$. For the VSD domain we would actually need 8589934592 transition probabilities for each $t \in T$.

In the graphical model, there are m state nodes, and with each node X_i is associated a transition probability assessment function γ_t^i for which at most $d \cdot d^{|\rho_s(X_i)|} \cdot d^{|\rho_t(X_i)|} \cdot 2^{|A_i|}$ probabilities have to be specified. If the numbers of synchronic and temporal predecessors are bounded by l_s and l_t, respectively, and the number of actions directly affecting X_i is bounded by q, for each $X_i \in X$, then the total number of parameters for the complete model is at most $m \cdot d \cdot d^{l_s} \cdot d^{l_t} \cdot 2^q = d^{1+l_s+l_t} \cdot 2^q$. Note that in the most extreme case, where $l_s = m-1$, $l_t = m$ and $q = k$, we obtain the number $m \cdot d^{2m} \cdot 2^k$. Fortunately, in most practical applications, we have both $l_s, l_t \ll m$ and $q \ll k$. For instance, in the VSD model, we have that $l_s = 2$, $l_t = 2$ and $q = 2$. The precise number of probabilities required for this model per moment $t \in T$ is 282, consisting of 136 stationary probabilities from synchronic dependencies and 146 time-dependent probabilities.

Summarising, the number of model parameters can be reduced by a factor of approximately $3 \cdot 10^7$ if we use a graphical model in the VSD domain. A further reduction in the amount of model parameters could be obtained by making assumptions on the interaction between stationary and time-dependent probabilities, and on how the effects of multiple actions may be combined.

6. CONCLUSIONS AND FURTHER RESEARCH

Clinical treatment planning requires the ability to predict the interplay between the natural history of disease and effects of intervening actions over time. Often, such predictions cannot be made with certainty, and trade-offs have to be made between the expected benefit of current and future decisions. We have discussed how treatment planning problems in paediatric cardiology can effectively be modelled using partially observable Markov decision processes (POMDPs). Furthermore, the dynamics of the treatment process under consideration were represented as temporal probabilistic networks in order to reduce the model complexity. The complexity of POMDP models is known to hamper their application in practice, and probabilistic networks have been suggested in the literature as a suitable means to reduce that complexity [4]. However, both theoretical and practical investigations supporting this claim were lacking so far. In this paper, we have demonstrated that for practical problems a significant reduction in the number of model parameters can be obtained with this representation.

A next step in the current line of research would be to perform a sensitivity analysis on the model given a choice of reward and utility functions, in order to assess the required precision in the estimates of various model parameters. Note that a rough probability estimate (i.e. having a broad plausible range) has to be improved only when the improved estimate may yield other optimal decision policies. Furthermore, the factorisation of the transition probability distributions of Eq. 1 can be used to enhance the efficiency of solution methods. Several forecasting algorithms for temporal belief networks have been described in the literature. Future investigations will have to point out whether these algorithms can be used in conjunction with dynamic programming techniques to construct optimal decision policies for real-life clinical treatment planning problems.

Acknowledgement. The investigations were (partly) supported by the Netherlands Computer Science Research Foundation with financial support from the Netherlands Organisation for Scientific Research.

REFERENCES

[1] F.J. Macartney, D.J. Spiegelhalter, and M.L. Rigby, "Medical management", in *Paediatric Cardiology*, R.H. Anderson, F.J. Macartney, E.A. Shinebourne, and M. Tynan, Eds., vol. 1, pp. 421–442. Churchill Livingstone, Edinburgh, 1987.

[2] K.J. Aström, "Optimal control of Markov decision processes with incomplete state estimation", *Journal of Mathematical Analysis and Applications*, vol. 10, pp. 174–205, 1965.

[3] W.S. Lovejoy, "A survey of algorithmic methods for partially observed Markov decision processes", *Annals of Operations Research*, vol. 28, pp. 47–66, 1991.

[4] C. Boutilier, T. Dean, and S. Hanks, "Planning under uncertainty: structural assumptions and computational leverage", in *New Directions in AI Planning*, M. Ghallab and A. Milani, Eds., pp. 157–171. IOS Press, Amsterdam, 1996.

[5] T. Dean and K. Kanazawa, "A model for reasoning about persistence and causation", *Computational Intelligence*, vol. 5, pp. 142–150, 1989.

[6] P. Dagum and A. Galper, "Forecasting sleep apnea with dynamic network models", in *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, 1993, pp. 64–71.

[7] S. Andreassen, R. Hovorka, J. Benn, K.G. Olesen, and E.R. Carson, "A model-based approach to insulin adjustment", in *Proceedings of the Third Conference on Artificial Intelligence in Medicine*, M. Stefanelli, A. Hasman, M. Fieschi, and J. Talmon, Eds., Berlin, 1991, vol. 44 of *Lecture Notes in Medical Informatics*, pp. 239–248, Springer Verlag.

[8] C. Berzuini, R. Bellazzi, and D. Spiegelhalter, "Bayesian networks applied to therapy monitoring", in *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, 1991, pp. 35–43.

[9] L. Kidd, D. Driscoll, and W. Gersony, "Second natural history study of congenital heart defects: results of treatment of patients with ventricular septal defects", *Circulation*, vol. 87, pp. 138–151, 1993.

[10] T.P. Graham and H.P. Gutgesell, "Ventricular septal defects", in *Heart Disease in Infants, Children, and Adolescents*, pp. 724–746. Williams & Wilkins, Baltimore, 1995.

[11] R.A. Howard, *Dynamic Programming and Markov Processes*, MIT Press, 1960.

[12] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley, New York, 1994.

[13] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, Palo Alto, 1988.

[14] N.B. Peek and J. Ottenkamp, "Developing a decision-theoretic network for a congenital heart disease", in *AIME '97: Proceedings of the Sixth Conference on Artificial Intelligence in Medicine Europe*, E. Keravnou, C. Garbay, R. Baud, and J. Wyatt, Eds. 1997, pp. 157–168, Springer Verlag.