MITLibraries
Document Services

## DISCLAIMER OF QUALITY

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available. If you are dissatisfied with this product and find it unusable, please contact Document Services as soon as possible.

Thank you.

Due to the poor quality of the original document, there is some spotting or background shading in this document.

# THE COMPLEXITY OF MIARKOV DECISION PROCESSES 

Christos H. Papadimitriou ${ }^{1}$ and John N. Tsitsiklis ${ }^{2}$


#### Abstract

We investigate the complexity of the classical problem of optimal policy computation in Markov decision processes. All three variants of the problem (finite horizon, infinite horizon discounted, and infinite horizon average cost) were known to be solvable in polynomial time by dynamic programming (finite horizon problems), linear programming, or successive approximation techniques (infinite horizon). We show that they are complete for $P$, and thercfore most likely cannot be solved by highly parallel algorithms. We also show that, in contrast, tite deterministic cases of all three problems can be solved very fast in parallel. The version with partially observed states is shown to be PSPACE-complete, and thus even less likely to be solved in polynomial time than the NP-complete problems; in fact, we show that, most likely, it is not possible to have an efficient on-line implementation (involving polynomial time on-line computations and memcory) of an optimal policy, even if an arbitrary amount of precomputation is allowed. Finally, the variant of the problem in which there are no obscrvations is shown to be NP-complete.


## 1. INTRODUCTION

In the past, Complexity Theory has been applied with considerable success to separate optimization problems with respect t.o their computational dificulty [GJ, PS]. Such results are valuable in that they succeed in formalizing the intuitive feeling by researchers in the area about the impossibility of certain approaches in solving some hard problents, and in some cases give a quite clear separation between the easy and the hard versions (special cases or gencralizations) of important problems. In most of this work, NP'completeness is the main notion of complexity used. The important distinction here is between problems that are in 1 , that is, can be solved in polynomial time (linear programming, max flow, miamum spanning tree) and those that are NP -complete, and therefore evidently cannot (such as integer programming and the traveling salesman problem).

Recent research in parallel algorithms [Col has suceceded in pointing out some important diferences between problems that are in $\Gamma$, that is, can be solved in polynomial lime. Some of these problems can be solved by algorithms that use cooperating processors so that the time requirements are reduced tremendously (such problems include, from the optimization point of view, minimum spanning tree problem, the shortest path problem, and others), whereas certain other problems (first of all linear programming, but also some special cases of it, such as maximum flow) do not appear to be susceptible'to such massive parallelization. NC is the class of problems that have algorithms using a polynomial number of processors, and time dotay which is polynomial in the logarithm of the input
${ }^{1}$ Departments of Computer Science and Oprerations Research, Stanford University. Rescarch supported by the National Science Foundation.

2 Department of Phactrical Engineering and Computer Science, M.I.T. Reseach supported by an IBiM facully development award.
size. The question then arises, is $\mathrm{NC}=\mathrm{P}$ ? In other words, is it true that all problems that can be solved satisfactorily by sequential computers, can also take the maximum possible advantage of parallel computers? Most researchers believe that the answer is negative.

The notions of reductions and completencss offer some important evidence here as well: Linear programming, maximum flow, and some other problems have been shown to be $P$-complete [DLS, GSS]. This means, intuitively, that such problems are as hard as any problem in $P$; thus, they can be massively parallelized only if $N C=P$, that is, if all problems in $P$ can, and thus there are no inherently sequential problems., Since this event is considered extremely unlikely, $P$-completencss is taken as evidence that a problem is not in NC, and thus cannot be satisfactorily parallelized.

The design of optimal strategies for Markov decision processes (say, for the discounted infinite horizon case) has been a very central and well-solved optimization problem. There are basically two approaches to its solution. First, the problem can be expressed as a linear program and thus solved (either by generic techniques, or specialized "policy improvement" routines [IIo]; this technique appears to be unpromising from the parallel algorithm point of view, since it involves linear programming, a problem that scems inherently sequential. Sccondly, there are iterative techniques that are known to converge fast on the optimal policy; however, such techniques are sequential also by nature. (Let us note in passing that the problem of deriving a "cleán" polynomial time algorithm for this problem, without using general linear programming or approximale techniques, is an important open question that has not been emphasized in the literature as much as it deserves.) An interesting question arises: Is the Markov decision problem in NC, or is it complete for $P$, and thus, in some sense, the use of lincar programming for its solution is inherent?

In this paper we show that computing the optimal policy of a Markov decision process is complete for $P$, and therefore most probably is inherently sequential. This is true also for the average cost case, as well as the finite horizon (non-stationary) case. However, we also show that the deterministic special cases of all three problems (hat is, in which all possible stochastic matrices have zero-one entries) is in NC.

Next, we address the question of the complexity of the partially observed generalization of this problem; this is an important problem arising in many applications in control theory and operations research [Be]. We show that the partially observed version of the finite horizon problem is PSPACE-complete. We do not cousider infinite horizon partially obscrved probiems because these are not combinatorial problems and do not seem to be exactly solvable by finite algorithms. Notice that only recently has there been work relating combinatorial optimization problems $[\mathrm{Or}, \mathrm{Pa} \mid]$ to this notion of complexity, which is stronger than $\mathrm{Nl}^{\prime}$-completeness (see Section 2 for detinitions). In fact, we show that, most likely, it is not possible to have an efficient on-line implementation (involving polynomial time on-line computations and mencory) of an optimal policy, even if an arbitrary amount of precomputation is allowed. Finally, we show that the same problem with no observations is NP-complete.

In the next Section we introduce the necessary concepts from Markov Decision Processes, and Complexity Theory. In Section 3 we present our results concerning the possibility of designing massively parallel algorithms for these problems. Finally, in Section 4 we prove our results on partially observable . Warkov processes.

## 2. DEFINITIONS

Complexity. For a general introduction to Computational Complexity see the bibliography in $\left[\mathrm{Pa}_{2}\right]$. Our model of (scquential) computation can be any one of the Turing machine, random access machine, etc. [ $A H U, \mathrm{LI}$ ]; all these models of computation are known to be equivalent within a polynomial. We denote by $P$ the class of all problems that can be solved in polynomial time (it is understood, in any such model of computation). As usual, by "problem" we mean an infinite set of instances, each of which is a "yes-no" question. Optimization problems can be transformed into such problems by asking, not for the optimal feasible solution, but simply whether the optimal cost is below a given bound.

NP is the class of problems that can be solved by a nondeterministic such machine in polynomial time; a nondeterministic machine essentially has the power of branching out and following both computation paths; this can be repeated in each of the paths, and so on, creating a potentially exponential set of computations.. (Equivalently, NP is the class of "yes-no" problems that have the "succinct certificate property", see [PS].)

Besides time complexity; we shall concern ourselves with the space complexity of algorithms. The space consumed by a computation is the number of memory cells necessary for the computation; notice that this does not count the space necessary for the input, so the space requirement of an algorithm may be less than linear (for example, logarithmic). An important class is PSPACE, the cliss of all problems that can be solved in polynomial space. It is truc that, if a computation can be carried out in time (or nondelerministic time) $T(n)$, then it requires only $T(n)$ space; it follows that $P$ and NP are subsets of PSPACE. If the space requirements are $T(n)$, then the time requirement cannot exceed $c^{T(r)}$, for some constant $c$.

We say that a problem $A$ is reducible to another $B$ if, for any instance $x$ of $A$, we can construct asing space $\log (|x|)$ an instance $\int(x)$ of $B$ such that the answer for $x$, considered as an instance of $A$, is the same as the answer of $f(x)$, considered as an instance of $B$ (here $x$ is a string encoding an instance of problem $\Lambda$, and $|x|$ is its length). Notice that we insist that the function $\int$ be computed in logarithmic space; this is stricter than the polynomial lime usually required. In practice this is not a major restriction as most known polynomial reductions can in fact be accomplished in logarithmic space. Also, it can be shown that log-space reduction is transitive. Let $\Lambda$ be a problem, and $C$ be a class of problems (such as $\mathrm{P}, \mathrm{N} \Gamma, \operatorname{lSP} \Lambda \mathrm{CE}$ ). We say that $\Lambda$ is complete for $\mathcal{C}$, or $\mathcal{C}$-complete if (a) $\Lambda$ is a member of $C$, and ( $b$ ) for every member $B$ of $C, B$ is reducible to $A$.

Nlthough there are many NP-complete problems in the literature, the notion of PSPACH-completencss is somewhat less well-known. A fundamental I'SI'ACF-complete problem is that of quantilied satisfiahifity ( $Q S A C$ ) |SiN|: We are given a quantified Boolean formula $\exists x_{1} \vee x_{2} \exists x_{3} \ldots \vee x_{n} F\left(x_{1}, \ldots, x_{n}\right)$, where $F$ is an ordinary Boolean formula in conjunctive normal form, with three literals per clanse. We are to determine whether this formula is true, that is, whether there exists a truth value for $x_{1}$ such that, for all truth values of $x_{2}$, etc. for all truth values of $x_{n}, F$ comes out true.

As mentioned above, $\operatorname{PSPACE}$ contains NP. In fact, ND can be thought of as the class of problems reducible to the special case of $\mathrm{OS} A \mathrm{~T}$, called $\mathrm{SA} \mathrm{A}^{\prime}$, in which all (quantifiers are existential (that is, there is no alternation of quatifiers). In fact, there is a number
of intermediate classes $\ddagger$ based on less restricted alternation; this is called the polynomial hierarchy [SM]. The next level above NP in this hicrarchy is the class $\Sigma_{2,}^{p}$ which can be thought of as all problems that reduce to the problem of telling whether a formula with one string of cxistential quantifiers, followed by a string of universal quantifiers, is true. An example of such a formula is $\exists x_{1} \exists x_{2} \ldots \exists x_{n} \forall y_{1} \forall y_{2} \ldots \forall y_{m} F\left(x_{1}, \ldots, y_{m}\right)$.

In the interest of formalizing parallel computation, we can define a parallel random access machine (PRAM), that is, a set of RAMs operating in unison, and communicating through a sequence of common registers, some of which initially hold the input [Co]. The measures of complexity here are the delay until the final answer is produced, and the number of processors that have to be involved. $\Lambda$ problem is considered to be solved in a satisfactory way by such a supercomputer if the delay is polynomial in the logarithm of the length of the input, and the number of processors involved are polynomial in the length of the input. The class of problems solvable under these terms is called NC. Obviously (by multiplying delay times number of processors) $N C$ is a subset of $P$. The great enigma for parallel computation, analogous to $\mathrm{P}=$ ? NP for sequential computation, is whether $\mathrm{NC}=\mathrm{P}$. That is, while $P=N P$ asks whether there are (reasonable) problems that are inherently exponential, $\mathrm{NC}=\mathrm{P}$ asks whether there are problems that are inherently sequential. The bets are that such problems do exist, but there is no proof. Now, if inherently sequential problems exist (equivalently, if $N C \neq P$ ) then the problems that are complete for $P$ will certainly be among them. Thus, $P$-completencss plays vis-a-vis the $N C=P$ problem the same role enjoyed by NP-completeness in relation to the $P=N P$ problem.

One basic P-complete problem is the circuit value problem (CVP). A circuit is a funite sequence of triples $C=\left(\left(a_{i}, b_{i}, c_{i}\right), i=1, \ldots, k\right)$. For each $i \leq k, a_{1}$ is one of the "operations" false, true, and, and or, and $b_{i}, c_{i}$ are non-negative integers smaller than i. If $a_{i}$ is cither false or true then the triple is called an input, and $b_{i}=c_{i}=0$. If $a_{i}$ is either and or or then the triple is called a gate and $b_{i}, c_{i}>0$. The value of a triple is defined recursively as follows: First, the value of an input (true, 0,0 ) is true, and that of (false, 0,0 ) is false. The value of a gate $\left(a_{i}, b_{i}, c_{i}\right)$ is the Boolean operation denoted by $a_{i}$ applied to the values of the $b_{i}$ th and $c_{i}$ th triples. The value of the circuil $C$ is the value of the last gate. Finally, the CVP is the following problem: Given a circuit $C$, is its value true?

Markov Decision Processes. In. a Markov decision process we are given a finite set $S$ of states, of which one, $\left.s_{0}\right)$ is the initial state. At each time $t=0,1, \ldots$ the current state is $s_{t} \in S$. For each state $s \in S$ we have a finile set $D$. of decisions. By making decision $i \in D_{n}$, at time $t$, we incur a cost $c(s, i, t)$, and a che next state $s^{\prime}$ has probability probability distribution given by $p\left(s, s^{\prime}, i, t\right)$. We say the process is stationary if $c$ and $p$ are independent of $t$. A policy $\delta$ is a mapping assigning to each lime $t \in\{0, t, 2, \ldots\}$ and state $s \in S$ a decision $\delta(s, t)$. A policy is stationary if $\delta(s, t)$ is independent of $t$. We shall consider three variants of the problem. First, in the finite horizon problem, we are given $a$ (non-stationary) Markov decision process and an inleger $T$. We wish to find a policy that minimizes the expectation of the cost $\sum_{t=0}^{T} c\left(s_{t}, \delta\left(s_{t}, l\right), t\right)$. The other variants that
$\dagger$ Of course, it should be borne in mind that, for all we know, PSPACE, ND, even $P$ might all coincide, in which case so would these "intermediate" classes.
we deal with have infinite horizon, and thus we shall only consider their stationary special cases (otherwise we nced the description of an infinite object, namely the parameters $c$ and $p$ for all times). In the discounted (infinite horizon) problem we are given a Markov decision process and a number $\beta \in(0,1)$, and we wish to minimize the expectation of the discounted cost $\sum_{t=0}^{\infty} c\left(s_{t}, \delta\left(s_{t}, t\right)\right) \beta^{t}$. Finally, in the average cost (infinite horizon) problem we are given a Markov decision process, and we wish to ininimize the expectation of the limit $\lim _{T \rightarrow \infty} \frac{\sum_{i=0}^{T} c\left(s_{2}, \delta\left(s_{t}, t\right)\right)}{T}$. It is well-known that in the two latter cases (those for which the optimal policy is a potentially infinite object), there exist stationary optimal policies, and therefore optimal policies admit a finite description. For all three cases of the problem, there are well-known polynomial-tinic algorithme.

## 3. P-COMPLETENESS AND PARALLEL ALGORITHMS

We show the following result:
Theorem 1. The Markov Decision Process problem is P -complete in all three cases (finite horizon, discounted, and average cost).

Proof: We shall reduce the CVP to the finite horizon version first. Given a circuit $C=\left(\left(a_{i}, b_{i}, c_{i}\right), i=1, \ldots, k\right)$, we construct a stationary Markov process $M=(S, c, p)$ as follows: $S$ has one state $i$ for each triple $\left(a_{i}, b_{i}, c_{i}\right)$ of $C$, plus an extra state $q$. If ( $a_{i}, b_{i}, c_{i}$ ) is an input, then the corresponding state $i$ has a single decision 0 , with $p(i, q, 0)=1$, and $\operatorname{cost} c(i, 0)=1$ if $a_{i}=$ false, and 0 otherwise. All other costs are 0 . From $q$ we only have one decision, with 0 cost, which has $p(q, q, 0)=1$. If $a_{i}$ is an orgate, there are two decisions 0,1 from state $i$, with zero cost, and $p\left(i, b_{i}, 0\right)=1, p\left(i, c_{i}, 1\right)=1$. That is, we decide whether the next state is going to be $b_{i}$ or $c_{i}$. If $a_{i}$ is an and gate, then there is one choice 0 in $l_{i}$, with $p\left(i, b_{i}, 0\right)=p\left(i, c_{i}, 0\right)=1 / 2$, that is, the next state can be either $b_{i}$ or $c_{i}$. We let the initial state to be $k$, the last gate, and $T$, the time horizon, to be equal to $k$, the size of the circuit.

We claim that the optimum expected cost is 0 or less if the value of $C$ was true. In proof, suppose that the expectation is indeed aero (it cannot be less). Then it follows that there are decisions so that the states with positive costs are impossible to reach. However, these states correspond to the false inputs, and thas these decisions are choices of a true gate among $b_{i}, c_{i}$ for each or gate $i$ of the circuil so that its overall value is true. Conversely, if the value is true, there must be a way to choose an input gate for each or gate so that the false inputs are not reachable, or, equivalently, the states with positive costs are not possible.

Fissentially the same construction works for the discounted ease; for the average eost problem, a modification is necessary: We must first make sure that all paths from the inputs to the last gate of $C$ have the same number of gates on them, and then have transitions from the states corresponding to the inputs not to a state $q$, but back to the initial stale.

We note that the above proof shows that even the stationary special case of the finite horizon problem is P-hard. It is not known whelher the stationary finite-horizon problem is in P, because of the following diflically: We could be given a stationary process with $n$ states, and a horizon $T$ up to $2^{n}$, and the input would still be of siac $O(n)$. Still, the
dynamic programming algorithm for this problem would take time proportional to $T n_{\text {, }}$ and thus exponential in $n$.

## Deterministic Problems

Consider now the special case of these problems, in which the only allowed values for $p$ are 0 and 1 ; we call this the deterministic case. We shall show below that the deterministic cases of the finite horizon (stationary and non-stationary), discounted, and average cost problem are in NC. Our approach is to look at these problems as variants of the graphtheoretic shortest-path problem [La]. This is done as follows: Given a deterministic Markov decision process, we construct a directed graph $G=(V, A)$, where the nodes are the states, and the arcs emanating from a node are the decisions available at this state. The node pointed by an arc is the state to which this decision is certain to lead (recall that there are only 0 or 1 probabilities). There is a weight $c(u, v)$ on each arc ( $u, v$ ), equal to the cost of the corresponding decision. We denote by $u_{0}$ the node corresponding to the initial state $s_{0}$. The particular variants of the problem are then equivalent to certain variants of the shortest path problem.

The Non-Stationary Finite Horizon Problem. The parallel algorithms that we describe in this Section employ a technique used in the past to yield fast parallel (or space-efficient) algorithms known as "path doubling" [Sa, SV]. The idea is, once we have computed all optimal policies between any two states, where each policy starts at time $t_{1}$ and ends at time $t_{2}$, and similarly between $t_{2}$ and $t_{3}$, to compute in one step all optimal policics between $t_{1}$ and $t_{2}$. We can think of this as nultiplying two $|V| \times|V|$ matrices $A\left(t_{1}, t_{2}\right)$ and $\Lambda\left(t_{2}, t_{2}\right)$, where the $(u, v)$ th entry of $\Lambda\left(t_{1}, t_{2}\right)$ is the cost of the optimal policy from state $u$ to $v$ between times $t_{1}$ and $t_{2}$. This "matrix multiplication" is the variant in which multiplication of reals is replaced by addilion, and addition of reals by the operation of taking the minimum. Using $n^{3}$ processors, we can "multiply" $n \times n$ matrices in $\log n$ parallel steps, by independently computing each of the $n^{2}$ inner products by $n$ processors in $\log n$ steps. This latter can be done by computing the $n$ terms of the inner product independently, and combining them in $\log n$ stages.

This approach immediately suggests an NC parallel algorithm for the finite horizon non-stationary problem: We start from the matrices $A(t, t+1), t=0, \ldots T-1$; the $(u, v)$ th entry of $A(t, t+1)$ equals to the cost of the decision leading at time $t$ from state $u$ to state $v$, if such a decision exists at time $t$, and equal to $\infty$ otherwisc. Then to compute the required $A(0, T)$ we multiply in $\log T$ stages these matrices. The total number of processors is $T n^{3}$, and the total paralle! time $\log 7 \log n$; since the size of the input for the non-stationary problem is $n^{2} T$, this is an NC algorithm.

Theorem 2. The finite-horizon, non-stationary deterministic problem is in NC. $\square$
Notice that this technique does not solve the stationary problem, whose input is of length $n^{2}+\log T$, and which will have to be attacked by a more sophisticated technique, explained later.

The Infinite Horizon Undiscounted Casc. It is easy to see that the infinite horizon average cost problem is equivalent to finding the cycle in the graph corresponding to the process that is reachable from $u_{1}$, and has the smallest average length of ares. In proof of
the equivalence, the limit of the average cost of an infinite path which starts at $u_{0}$, reaches: this cycle, and then follows this cycle for ever, equals to the optimum average cost, and this cannot be improved.

To make sure that we do not consider solutions that are not reachable from $u_{0}$, we first compute in $\log n$ parallel time [SV] the nodes reachable from $u_{0}$, and in the sequel: consider only these. The.cycle with the shortest average cost can be found by computing; in parallel, for each $k=1, \ldots, n$ the shortest cycle of length $k$, and comparing the results, each divided by $k$. To compute the shortest cycle of length $k$, we essentially have to compute the $k$ th power of matrix $A$, whose $(u, v)$ th entry is equal to the cost of the decision leading from state $u$ to state $v$, if such a decision exists, and $\infty$ otherwise. This. is done with $n^{2} k$ processors in $\log k \log n$ parallel steps. The total time is $\log ^{2} n+2 \log n$, using $n^{4}$ processors.
Theorem 3. The infinite-horizon, avarage cost deterministic problem is in NC.
The Infinite Horizon, Discounted Case. Define a sigma in a directed graph to be a path of the form ( $u_{0}, \ldots, u_{k}, v_{1}, \ldots, v_{\ell}, v_{1}$ ), where all nodes indicated are distinct. In other words, a sigmia is a path from $u_{0}$ until the first repetition of a node. The discounted cost of a sigma $P=\left(u_{0}, \ldots, u_{k}, v_{1}, \ldots, v_{\ell}, v_{1}\right)$ is defined to be $c(P)=\sum_{i=0}^{k} c\left(v_{j}, v_{j+1}\right) \beta^{i}+$ $\frac{3^{k+1}}{\left(1-\beta^{\prime}\right)} \sum_{j=1}^{\ell} c\left(u_{i}, u_{i+1( }(\right.$ mocl $\left.\ell)\right) \beta^{j}$. That is, the discounted cost of a sigma coincides with the discounted cost of an infinite path that follows the sigma and repeats the cycle for ever. It follows from the fact that the (general) discounted problem has a stationary optimal policy, that the cost of chis optimal policy is the optimum discounted cost of a sigma in the corresponding directed graph.

We can compute the optimum sigma as follows: First, we compute the shortest discounted path of length $j=0,1 \ldots, n$ among any pair of nodes by "multiplying" the matrices $A, \beta A, \ldots, \beta^{j-1} A$, where $A$ is the matrix defined in the paragraph before Theorem 3. This can be done in $\log ^{2} n$ steps by using $n^{4}$ processors. Let $B_{1}, \ldots, B_{n}$ be the resulting products; the $(u, v)$ th entry of $B_{j}$ is the length of the shortest path with $j$ ares from $u$ to $v$. Once this is done, for each node $u$ and for each $k, \ell=0,1 \ldots, n$ we compule $\lambda+\frac{\beta^{k+1}}{1-\beta^{7}} \mu$, where $\lambda$ is the $\left(u_{( }, u\right)$ th entry of $B_{k}$, and $\mu$ is the $u, u$ entry of $B_{l}$. We pick the best result. This requires a total of $n^{4}$ processors, and $\log ^{2} n+3 \log n$ parallel steps.

Theorem 4. The infinite-horizon, discounted deterministic problem is in NC. $\square$
The Finite Horizon, Stationary Case. Given a stationary deterministic process, the stationary linite horizon problem with horizon $T$ is equivalent to finding the shortest path with $T$ ares in the corresponding graph, starting from the node $u_{1}$. The algorithm should run in a mumber of paralled steps that is polynomial to log lor $T$; therefore the "pathdoubling" lechnique would not give the desired result. In the sequel we assume that T' $>n^{2}$, otherwise the previous technique applies.

Without loss of generality, assume that the are lengths are such that no ties in the lengths of paths are possible (this can be achieved by perturbing the lengths). Consider the shortest path out of $u_{0}$ wilh 7 ares. We can consider this path as a simple path out of $u_{0}$, plus several simple cycles. We first need to show that we can assume that only one simple cycle is repeated more than $n$ times, namely the one that has the shortest average
length of its arcs. In proof, consider two cycles of length $m_{1}$ and $m_{2}$, repeated $n_{1}$ and $n_{2}$ times, respectively. If $n_{1}, n_{2} \geq n$, then we can repeat the cycle with smaller average are length, say the first, $m_{2}$ times less, and repeat the second $m_{1}$ times more, to obtain another path of sinaller length. Thercfore, the shortest path of length $T$ has the following structure: It consists of a path with length $\ell<n^{2}$, plus a simple cycle repeated many times to fill the required number of arcs. Therefore, for each value of $\varepsilon<n^{2}$ and each node $u$ and each possible number of arcs in the cycle $k$ that divides $T-\ell$ we do, in parallel, the following: We compute the shortest path of length $\ell$ from $u_{0}$ through $u$, the shortest cycle of length $k$ through $u$, and the cost of adding $(T-\ell) / k$ copies of the cycle to the path. Of the resulting combinations, we pick the cheapest.

Theorem 5. The finite-horizon, stationary deterministic problem is in NC.

## 4. THE PARTIALLY OBSERVED PROBLEM

We can define an important gencralization of Markov decision processes if we assume that our policy must arrive at decisions with only partial information on the current state. In other words, we have a partition $\Pi=\left\{z_{1}, \ldots, z_{p}\right\}$ of $S$, where the $z_{i}$ 's are disjoint subsets; of $S$ that exhaust $S$, and at any time we only know the particular set in $\Pi$ to which the current state belongs (naturally, we also remember the previous such sets) $\dagger$. Each set $z \in \Pi$ has a set $D_{z}$ of decisions associated with it, and each decision $i \in D_{z}$ has a cost $c(z, i)$ - and a probability distribution $p\left(s, s^{\prime}, i\right)$ of the next state $s^{\prime}$, given the current state $s \in z$. The initial state is known as before to be $s_{0}$. A policy is now a mapping from scquences of oljservations $z_{1}, \ldots, z_{t}$ to decisions in $D_{z_{t}}$. This problem is often treated by converting it to an equivalent perfectly observed problem by redefining the state to be the vector $p$ whose $i$-th component is the conditional probability that the state of the original problem equals $i[\mathrm{Bc}]$. In this formulation the state space is infinite. Nevertheless, it is known [ SS ] that the optimal cost-to-go function is of the form $J(p, t)=\min _{k \in K_{t}}\left[\Lambda_{k}+B_{k} p\right]$, where $A_{k}$ is a scalar and $B_{k}$ is a vector of suitable dimension and $K_{t}$ is a finite index set. This allows the solution in finite time of the problem even though the state space is infinite. Nevertheless, the cardinality of the index set $K_{t}$ may be an exponential function of the time horizon and, therefore, this procedure is of limited practical use. The following result shows that most likely this is a generic property of the probiem we are dealing with rather than a defect of the particular reformulation of the problem. Notice also that in the infinite horizon limit the index set $K_{t}$ will be infinite, in general, which makes fairly doubtful whether there exists a finite algorithm for the infinite horizon problem. For this reason, we only consider the problem in which we are asked to minimize the expected undicounted cost over a finite Lime horizon 7 .

Theorem 6. The partially observed problem is I'SIPACE-complete, even if the process is stationary and $T$ is restricted to be smaller than $|S|$.
$\dagger$ In the traditional formulation of the partially observed problem [Be], the observation at time $t$ is a random variable, for which we know its conditional distribution given the current state and time. Here we essentially deal with the special case in which the observation is a deterministic function of the curent state. Clearly, the general case cannot be any casier.

Proof: We first sketch a proof that the problem is in PSPACE. The construction is a bit tedious, but quite reminiscent of polynomial-space algorithms for other similar problems (see for example $[\mathrm{Pal}]$ ). We can think of the possible outcomes of the process within the finite horizon $T$ as a tree of depth $T$, which has internal nodes both for decisions and transitions. The leaves of this tree can be evaluated to determine the total contribution to the expected cost of this path. To determine whether the optimal policy has expected cost less than some number $B$, we need to traverse this tree, making nondeterministic decisions at the decision nodes, iterating over all transitions, and taking care to make the same decision every time the history of obscrvations is identical. This latter can be achieved by organizing the search in such a way that nodes at the same level of the tree are divided into intervals from left to right, corresponding to distinct observation sequences. Decisions from nodes in the same interval must be the same. That the problem is in PSPACE follows now from the fact that PSPACE is robust under nondeterminism.

To show PSPACE-completeness, we shall reduce QSAT to this problem. Starting from any quantified formula $\exists x_{1} \forall x_{2} \ldots \forall x_{n} F\left(x_{1}, \ldots, x_{n}\right)$, with $n$ variables (existential and universal) and $m$ clauses $C_{1}, \ldots, C_{m n}$, we construct a partially observed stationary Markov process and a horizon $T<|S|$, such that its optimum policy has cost 0 or less if and only if the formula is true.

We first describe $S$. Besides $s_{0}$, the initial state, $S$ contains six states $A_{i j}, A_{i j}^{\prime}, T_{i j}, T_{i j}^{\prime}, F_{i j}, F_{i j}^{\prime}{ }^{\text {M }}$ for each clause $C_{i}$ and variable $x_{j}$. There are also $2 m$ states $A_{i, n+1}, A_{i, n+1}^{\prime}$. We next describe the partition $\Pi$. The initial state $s_{0}$ is in a sct by itself. For cach variable $j$, the states $A_{i j}, i=1, \ldots, m$ and the states $A_{i j}^{\prime}, i=1, \ldots, m$ form a set called $A_{j}$, the states $T_{i j}, i=1, \ldots, m$ form a set called $T_{j}$, the states $T_{i j}^{\prime}, i=1, \ldots, m$ form a set called $T_{j}^{\prime}$ and so on, up to $F_{j}^{\prime}$. The states $A_{i, n+1}$ and $A_{i, n+1}^{\prime}$ are each on a different set.

Next we shall describe the decisions, the probabilities, and the costs. All decisions except those out of the set $\Lambda_{n+1}^{\prime}$ have zero cost. At $s_{0}$ there is only one decision, leading to the states $A_{i 1}, i=1, \ldots, m$ with equal probability. If $x_{j}$ is an existential variable, there are two decisions out of the set $\Lambda_{j}$, leading certainly from $A_{i j}$ to $T_{i j}$ and $F_{i j}$, respectively; similarly there are two decisions ont of the sel $\Lambda_{j}^{\prime}$, leading with certainty from $\Lambda_{i j}^{\prime}$ to $J_{i j}^{\prime \prime}$ and $F_{i j}^{\prime \prime}$, respectively. If $x_{j}$ is a universal variable, there is one decision out of the set $\Lambda_{j}$, leading with equal probability from $\Lambda_{i j}$ to $T_{i j}$ and $F_{i j}$; similarly there is one decision out of the set $A_{j}^{\prime}$, leading with equal probability from $\Lambda_{i j}^{\prime}$ to $T_{i j}^{\prime}$ and $F_{i j}^{\prime}$. From the $T_{j}, F_{j}, T_{j}^{\prime}$, and $F_{j}^{\prime}$ sets there is only one decision, which leads with certainty Prom $T_{i j}, F_{i j}, T_{i j}^{\prime}$, and $F_{i j}^{\prime}$ to (respectively) $A_{i, j+1}, A_{i, j+1}, \Lambda_{i, j+1}^{\prime}$, and $\Lambda_{i, j+1}^{\prime}$, with t.wo exceptions: If $x_{j}$ appears positively in $C_{i}$, the transition from $T_{i j}^{\prime}$ is to $\Lambda_{i . j+1}$ instead of $\Lambda_{i . j+1}^{\prime}$; and if $x$, apears negatively, the transition from $P_{i, j}^{\prime}$ is .on $\Lambda_{i, j+1}$.

Finally, out of $A_{i, n+1}$ and $A_{i, n+1}^{\prime}$ there is one decision, leading to the same state with ecrtainty, only that the decision out of $\Lambda_{i . n+1}^{\prime}$ incurs a cost of 1 (the only one with nonzero cost). This completes the construction of the process; the horizon $T$ is delined to be $2 m+2$ (just enough time for the process to reach one of $\Lambda_{i, n+1}$ or $\Lambda_{i, n+1}^{\prime}$ ).

We claim that there exists a policy with expected cost zero iff the formula is true. Suppose such a policy exists. Recall that the lransition from the initial state can be to any state $A_{i 1}^{\prime}$ of $A_{1}^{\prime}$; we think that the process "chooses" a dause (\%. Once this has happened, the process remains at states subseripted by $i$ forever, without ever observing
the real value of $i$. It follows that the policy has zero expected cost for all such initial choices. We next claim that the policy must guarantee that the process ends up in $A_{i, n+1}$. If not, that is, if for some choices of decisions for the universal variables the process ends up in $A_{i, n+1}^{\prime}$, then this contributes an expected cost of at least $2^{-\pi} / \mathrm{m}$, which is absurd. It follows that the policy must, based on the previous observations on the transitions at the sets corresponding to previous universal variables, pick at the existential variables those decisions that correspond to a truth assignment which satisfies the clause. .Since all clauses must be satisfied, the formula was true.

Converscly, if the formula were true, there is a policy for setting the existential variables, based on the values of previous universal ones, so that all clauses are satisfied. This, however, can be translated into a policy for choosing the corresponding decisions at the sets corresponding to existential variables so that, for all choices of clauses, the state $A_{i, n+1}$ is reached. A policy with zcro cost is thus possible. The proof is complete. $\square$

Theorem 6 says that, unless $P=P S P A C E$ (an event even less likely than $P=N P$ ), there is no polynomial-time algorithm for telling whether a cost can be achieved in a given partially observable process. However, there is a more practical problem, the complexity of which is not settled by this result. It would still be satisfactory if we could analyze such a process (perhaps in exponential time or worse), and, based on the analysis, design a practically realizable controller (polynomial-time algorithm) for driving the process online optimally. This algorithm would presumably consult some polynomially large data structure, constructed during the analysis phase. We point out below that this is very unlikely:

Corollary 1. Unless PSPACE $=\Sigma_{2}^{p}$, it is impossible to have an algorithm $A$ and a mapping $\mu$ (of arbitrarily high complexity) from a partially observed Markov process $M$ to a string $\mu(M)$ of length polynomial in the description of the process, such that the algorithm, with the string and the observations as input, computes the optimum decision for the process in polynomial time.

Sketch: If such an $A$ and $\mu$ existed, then we could express the question of whether $M$ has zero expected cost (and thus, by the previous prool, whether an arbitrary quantified Boolean formula is true) as follows: "There exists a string $\mu(M)$ such that, for all possible transitions, the decisions of algorithm $A$ lead to a zero cost." This last sentence, however, can be rewritten as a Boolean formula with two alternations of quantifiers.

Finally, considering the special case of the partially observed problem, in which the partition $I I=\{S\}$ (i.e., the case of mobserved processes or, equivalently, an open-loop control) we note that a simplification of the proof of Theorem $G$ establishes that this problem is NI'complete:

Corollary 2. Deciding whether the optimal policy in an unobserved Markov decision process has expected cost equal to zero is an NP-complete problem. $\square$

Bibliopgraphy
|AllU| A.V. Aho, J.E. Hoperoft, J.D. Ullman The Design and Analysis of Computer Algorithms, Addison, 1974.
[Be] D. 「' Bertsekas, Dynamic Programming and Stochastic Control, Academic Press, 1976.
[Co] S.A. Cook "Towards a Complexity Theory of Synchronous Parallel Computation" Enseign. Math. 2, 27, pp. 99-124, 1981.
[DLR]D. Dobkin, R.J. Lipton, S. Reiss "Linear Programming is Log-Space Hard for P", Inf. Proc. Letters 8, 96-97, 1979.
[GSS] L.M. Goldschlager, R.A. Shaw, J. Staples "The Maxinum Flow Problem is Log Space Complete for P", Theor. Computer Science 21, pp. 105-111, 1982.
[GJ] M.R. Garey, D.S. Johnson Compuiers and Intractability: A Guide to the Theory of NP-completeness, Freemen, 1979.
[Ho] R.A. Howard Dynamic Programming and Markov Processes, MIT Press, 1960.
[La] E.L. :Lawler Combinatorial Optimization: Networks and Matroids, Holt, Rinehart and Winston, 1977.
$[L P]$ H.R. Lewis, C.H. Papadimitriou Elements of the Theory of Computation, PrenticeHall, 1981.
[Or] J. Orlin "The Complexity of Dynamic Languages and Dynamic Optimization Problems", Proc. 13th STOC, pp. 218-227, 1981.
[Pa2] C.H. Papadimitriou "Computational Complexity", in Combinatorial Optimization: Annotated Bibliographies, M. O' hEigeartaigh, J.K Lenstra, A.H.G. Rinooy Kan (eds.), pp. 39-51, 1985.
[Pa1] C.H. Papadimitriou "Games Against Nature", J. Comp. and Systems Sci., to appear, 1985.
[PS] C.H. Papadimitriou, K. Steiglitz Combinatorial Optimization: Algorithms and Complexity, Prentice-Hall, 1982.
[Sa] W.J. Savitch "Relationships Between Nondeterministic and Deterministic Tape Complcxities", J. Comp. and Sys. Sci. 4, pp. 177-192, 1970.
[SS] R.D. Smallwood, E.J. Sondik, "The Optimal Control of Partially Observable Markov Processes over a Finite Horizon", Operations Rescarch, 11, pp. 1071-1088, 1973.
[SM] L.J. Stockmeyer, A.R. Meyer "Word Problems Requiring Exponential Space", Proc. 5th STOC, 1973.
[SV] Y. Shiloach, U. Vishkin " $\operatorname{An} O(\log n)$ Paralle! Conneclivity Algorithm", J. Algorithms, 3 рр. 57-67, 1982.

