

An Ant System Approach to Markov Decision Processes

Hyeong Soo Chang, Walter J. Gutjahr, Jihoon Yang, and Sungyong Park

Abstract—In this paper, we develop an ant-system based algorithm for approximately solving large Markov Decision Process (MDP) problems for infinite horizon discounted cost criterion, extending the applicability of the ant-system metaheuristic into stochastic sequential decision making problems. The algorithm inherits the spirit of the well-known policy iteration algorithm with an adaptation of the ant system into MDP settings with some modifications and extensions, while preserving the probabilistic convergence property of the ant system.

I. INTRODUCTION

A novel “general purpose” or *metaheuristic* algorithm based on the foraging behavior of ant colonies has been applied, with suitable modifications, to various *deterministic* combinatorial optimization problems to break *the curse of dimensionality* in solving them and it has been reported that the results are successful in many problems (see, e.g., [10] [11] for surveys). Recently, it has been shown in [16] that the solutions generated at each iteration of an ant-system based algorithm converge to the unique optimal solution of a given combinatorial optimization problem, under certain conditions, with a probability that can be made arbitrarily close to 1.

In this paper, we develop an ant-system based algorithm for approximately solving Markov Decision Processes (MDPs) [21] for infinite horizon discounted cost criterion based on the ant-system framework of [16], extending the applicability of the ant system metaheuristic into *stochastic* optimization problems. Many stochastic control optimization problems formulated by MDPs experience the curse of dimensionality, making it very impractical to apply the well-known exact algorithms, e.g., Value Iteration (VI) and Policy Iteration (PI) [21], to solve the MDPs. There are numerous approximation-based schemes to get away with the dimensionality problem via various techniques, e.g., structural analysis, aggregation, sampling, feature extraction, etc. See, e.g., [3], [21], [22]. However, to the best of the authors’ knowledge, there is few algorithm that incorporates the success of “evolutionary computation” into solving MDPs except the recent works of Chang et al. [5] [4] based on genetic algorithm (GA) [15] and simulated annealing [17].

Ant System was developed by Dorigo and other authors (see [12]) and has been generalized later on to the Ant Colony Optimization (ACO) metaheuristic approach [10]

H. S. Chang, J. Yang, and S. Park are with Department of Computer Science and Engineering at Sogang University, Seoul, Korea and can be reached by e-mail at {hschang, jhyang, parksy}@sogang.ac.kr.

W. J. Gutjahr is with Department of Statistics and Decision Support Systems, University of Vienna, Vienna, Austria and can be reached by email at walter.gutjahr@univie.ac.at.

which has found numerous applications in diverse fields in the meantime. It is a constructive and repetitive heuristic technique where the constructive part is similar to greedy heuristics, and the repetitive part draws from concepts in Reinforcement Learning [22]. The best way of explaining the basic ideas of ACO is by reference to a biological metaphor which also gave the approach its name: the search behavior of biological ant colonies. In nature, ants move around in their environment in a rather random way, but they have a certain tendency to follow the walk of other ants. They can recognize these walks because, while moving, each ant leaves a chemical substance called *pheromone* on the ground. Sensing pheromone on a path increases the probability of an ant to follow it, which further reinforces this path. This mechanism has the effect that *short* paths between a starting point and a goal point are favored, leading to a kind of heuristic optimization behavior.

In ACO, the described principle is exploited for optimizing an arbitrary objective function of a combinatorial problem by simulating the walks of conceptual ants and by doing the reinforcement of good walks based on a post-hoc evaluation of the objective. In the ant system [16] that we consider in this paper, a directed graph (called “Construction Graph”) is used as the basic structure on which the walks take place. The feasible solutions of the combinatorial optimization problem are encoded as paths on this graph. In each iteration, several ants are first positioned at the same node of the graph. Then, in successive steps, each ant chooses randomly a successor node of the node where it is currently located. The probabilities for the successor nodes are computed from “pheromone values” assigned to the arcs of the graph (and additionally from so-called “visibility values”, which are pre-evaluations of the suitability of possible continuations). After all ants have finished their walks, the walks are decoded as problem solutions and evaluated. The arcs on “good” walks are then reinforced by pheromone increments. This process is iterated. We will express the procedure in more formal terms further below (see, Section III).

The ANT-PI algorithm proposed in this paper inherits the spirit of the PI algorithm for solving MDPs exactly. Loosely speaking, the basic algorithmic procedure is as follows: each ant generates its candidate policy (traversing a directed graph) and evaluates its policy (“policy evaluation” step in PI). The (estimate) quality of each ant’s candidate policy is merged to deposit a new pheromone, which in turn is used for each ant to generate a new candidate policy. During the merge step of the new pheromone generation, an *elite* policy is generated via a method called “policy switching” [6], where the elite policy is a policy that

improves the performances of all of the candidate policies generated so far by ants, based on an idea similar to the “elitist strategy” concepts [12], and the quality of the elite policy is used for the new pheromone generation.

Contrary to PI, by directly manipulating the policies, ANT-PI eliminates the operation of minimization over the entire action space in the “policy improvement” step of PI and a certain monotonicity property is preserved over the elite policies. Therefore, the proposed algorithms would be effective in the domains where the action state space is finite but extremely large (see, Section IV). If a local search is used, it cannot guarantee that a global minimum has been found. The main contribution of the present paper is the use of a (random) evolutionary search algorithm in the context of MDPs with a probabilistic convergence guarantee. ANT-PI preserves a certain monotonicity property via elite policies and converges to an optimal policy with a probability that can be made arbitrarily close to 1.

This paper is organized as follows. We start with some necessary background on MDPs in Section II. In Section III, we formally describe ANT-PI with detailed discussion. In Section IV, we discuss some complexity issue of ANT-PI compared with PI. We then conclude the present paper with some remarks in Section V.

II. MARKOV DECISION PROCESS

We briefly provide the formal model of Markov decision processes. For a substantial discussion, see, e.g., [21].

An MDP M is 4-tuple (X, A, P, C) , where X is a finite set of states in the system and A is a set of actions. At state x , not all actions in A can be taken. We let $A(x)$ denote the finite set of *admissible* actions in A at state $x \in X$. P is a state transition function. We denote the probability that a state $x \in X$ makes a transition to another state $y \in X$ by taking an action $a \in A(x)$ as $P(y|x, a)$. C is a bounded cost function that maps $X \times A(X)$ to a real number in \mathcal{R} .

We define a policy π as a mapping from X to $A(X)$ and let Π be the set of all possible policies and we define *the value of following a policy* $\pi \in \Pi$ with initial state $x \in X$ as

$$V^\pi(x) = E \left[\sum_{t=0}^{\infty} \gamma^t C(X_t, \pi(X_t)) \middle| X_0 = x \right]$$

where X_t is a random variable that denotes the state at time t and $\gamma \in (0, 1)$ is a fixed discount factor.

Given a fixed initial state probability distribution δ defined over X , we define *the average value of π for δ* :

$$J_\delta^\pi = \sum_{x \in X} V^\pi(x) \delta(x)$$

and *the average optimal value with δ* :

$$J_\delta^* = \min_{\pi \in \Pi} J_\delta^\pi.$$

The goal is to find an optimal policy $\pi \in \Pi$ that achieves J_δ^* .

The following result, Bellman’s optimality principle, is well-known in the literature [21].

Theorem 1: Let $V^*(x) = \min_{\pi \in \Pi} V^\pi(x)$, $x \in X$. For all $x \in X$,

$$V^*(x) = \min_{a \in A(x)} \left\{ C(x, a) + \gamma \sum_{y \in X} P(y|x, a) V^*(y) \right\}$$

and $V^*(x)$, $x \in X$ is unique and a policy π^* defined as

$$\pi^*(x) \in \arg \min_{a \in A(x)} \left\{ C(x, a) + \gamma \sum_{y \in X} P(y|x, a) V^*(y) \right\}$$

for $x \in X$ satisfies that $V^{\pi^*}(x) = V^*(x)$, $x \in X$, achieving J_δ^* for any δ .

There exist two well-known algorithms to compute π^* or the optimal value function V^* . VI starts with an arbitrarily chosen (bounded) function defined over X and updates the function into a new function defined over X . PI starts with arbitrarily chosen initial policy and generates a new policy which improves the old policy. PI computes π^* in a finite number of steps because there are a finite number of policies in Π and PI preserves the monotonicity in terms of the policy performance. The running time-complexity of VI is polynomial in $|X|$, $|A|$, and $1/(1 - \gamma)$ and in particular one iteration takes $O(|X|^2|A|)$ time. For PI, performing the policy improvement step takes $O(|X|^2|A|)$ time. See [18] for a detailed discussion, including the state and action space dependent time-complexity of the linear programming approach for solving MDPs. Therefore, applying the exact methods for solving MDPs is very difficult if the state and/or the action space are large, which is true for many interesting problems. We will discuss complexity issues further in Section IV after we present our algorithm.

III. ANT-PI ALGORITHM

A. Algorithm description

A high-level description of ANT-PI is shown below, where some steps are described at a conceptual level, with details provided in the following subsections.

Ant System Policy Iteration (ANT-PI)

- **Initialization:**

Select the number of ants n . Set $k = 1$.

Select $\mu \in (0, 1]$ and a positive function φ and an arbitrary $\pi_0^e \in \Pi$ and set $J_\delta^{\pi_0^e} = \infty$.

Initialize the pheromone for each ant: set $\tau_0(x, a) = 1/|A(x)|$ for all $x \in X$ and $a \in A(x)$. Initialize the policy generation function: set $\rho_0(x, a) = 1/|A(x)|$ for all $x \in X$ and $a \in A(x)$.

- **Repeat:**

- **Policy Generation (Walks of Ants):**

- * Generate $\Lambda_k = \{\pi_k^i, i = 1, \dots, n\}$, where π_k^i is generated using ρ_{k-1} .

- * Obtain $J_\delta^{\pi_k^i}$ for each $\pi \in \Lambda_k$.

– **Pheromone Update:**

- * For each $i = 1, \dots, n$, for each $x \in X$ and $a \in A(x)$,

$$\tau_k^i(x, a) = \begin{cases} \varphi(J_\delta^{\pi_k^i}) & \text{if } \pi_k^i(x) = a \\ 0 & \text{otherwise.} \end{cases}$$

- * For each $x \in X$ and $a \in A(x)$,

$$\tau_k(x, a) \leftarrow f(\tau_{k-1}(x, a), \tau_k(x, a), \mu) \quad (1)$$

– **Policy Generation Function Update:**

- * For each $x \in X$ and $a \in A(x)$,

$$\rho_k(x, a) = \frac{\tau_k(x, a)}{\sum_{a' \in A(x)} \tau_k(x, a')}.$$

– **Elite Policy Generation (Policy Switching):**

- * Generate the elite policy of Λ_k defined as

$$\pi_k^e(x) \in \left\{ \arg \min_{\pi \in \Lambda_k \cup \{\pi_{k-1}^e\}} (V^\pi(x))(x) \right\}, x \in X. \quad (2)$$

- $k \leftarrow k + 1$

B. Policy generation

We can view the set Π of all possible stationary policies as a “construction graph” in the ant system [16]. The construction graph is a pair (G, Ω) , where $G = (V, E)$ is a directed graph such that there is a unique starting vertex or node from which each ant begins traversing the graph. Each directed *walk* of an ant in G contains each node of G at most once and the last node on the walk has no successor node in G , that is not already contained in the walk. The function Ω maps the set of ants’ feasible walks to the set of feasible solutions of a given optimization problem. That is, the construction graph specifies a particular encoding of the feasible solutions as walks of ants. The objective function value of the walk is set equal to the objective function value of the corresponding feasible solution of the original problem.

In our contexts, we arrange the states in X in an arbitrary order \mathcal{O} . Each state $x \in X$ is a node $v \in V$ in the construction graph and each arc e corresponds to a pair (x, a) of a state $x \in X$ and an admissible action $a \in A(x)$ at the state and the direction of the arc reaches from the current state x to the next state in the order \mathcal{O} . A particular ant traverses all the states in the construction graph following the order \mathcal{O} from the first state of \mathcal{O} . When it moves from a state x to another state y , it traverses randomly the arc $e = (x, a)$ with the transition probability specified by $\rho_{k-1}(x, a)$, $a \in A(x)$, where k is the particular iteration step of ANT-PI and it is initialized with uniform distribution such that $\rho_0(x, a) = 1/|A(x)|$ for all $x \in X$ and $a \in A(x)$. It can be directly seen that once an ant finishes the tour, it has generated a stationary policy $\pi \in \Pi$ at random.

After random stationary policies π_k^1, \dots, π_k^n are generated from the walks of the ants, each policy is evaluated obtaining the value of following the policy. For each policy

$\pi \in \Lambda_k = \{\pi_k^1, \dots, \pi_k^n\}$, the following set of linear equations are solved for $V^\pi(x), x \in X$:

$$V^\pi(x) = C(x, a) + \gamma \sum_{y \in X} P(y|x, a) V^\pi(y), x \in X,$$

obtaining J_δ^π by averaging $V^\pi(x)$ over δ .

Once this “forward” phase (walking from the starting state to the ending state in \mathcal{O}) is finished, the next phase is “backward”. Each ant traverses back along the path it travelled and deposits its pheromone on the path. Depositing its pheromone along the path is like expressing its opinion on the policy he generated to the other ants. In the next subsection, we study how this pheromone is generated.

C. Pheromone and policy generation function update

While each ant backtracks the path he traversed over the construction graph at iteration k , each ant $i = 1, \dots, n$ modifies the pheromone τ_k^i . The pheromone is initialized with uniform distribution such that $\tau_0(x, a) = \frac{1}{|A(x)|}$ for all $x \in X$ and $a \in A(x)$. At iteration k , for each $i = 1, \dots, n$, for each $x \in X$ and $a \in A(x)$,

$$\tau_k^i(x, a) = \begin{cases} \varphi(J_\delta^{\pi_k^i}) & \text{if } \pi_k^i(x) = a \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where φ is a non-increasing function which may depend on the walks of the ants at the iterations $1, \dots, k-1$. (We defer the discussion on the function φ and the condition for the use of φ to the next subsection.) We then *merge* each ant’s pheromone by the function f in Equation (1) in Section III-A given as follows: for each $x \in X$ and $a \in A(x)$, if

$$\sum_{x' \in X, a' \in A(x')} \sum_{i=1}^n \tau_k^i(x', a') = 0$$

then, $f(\tau_{k-1}(x, a), \tau_k(x, a), \mu) = \tau_{k-1}(x, a)$. Otherwise,

$$f(\tau_{k-1}(x, a), \tau_k(x, a), \mu) = (1 - \mu)\tau_{k-1}(x, a) + \mu \cdot \frac{\sum_{i=1}^n \tau_k^i(x, a)}{\sum_{x' \in X, a' \in A(x')} \sum_{i=1}^n \tau_k^i(x', a')}.$$

The constant $\mu \in (0, 1]$ is called the *evaporation factor* (see, e.g., [12]). If no walk is rewarded ($\sum_{x' \in X, a' \in A(x')} \sum_{i=1}^n \tau_k^i(x', a') = 0$), the pheromones for all pairs of $(x, a), x \in X, a \in A(x)$ do not change. Otherwise, only an amount of $1 - \mu$ of pheromone remains on the graph after iteration k . The remaining amount of μ is the “budget” for rewarding the walks traversed at iteration k according to their respective average value of following their (randomly generated) policies. Each ant i can be thought as depositing the amount of $\tau_k^i(x, a)$ to express the degree of “goodness” on its walk at iteration k . Then, proportionally to the reported values $\tau_k^i(x, a)$, the actual pheromone is updated within the budget of μ .

By the update rule, the pheromone on “promising” arcs of the construction graph is increased, such that those arcs are traversed more often by the ants in the future. At the

extreme case of $\mu = 0$, the system ignores the impact of the value function on the walks of the ants, which is excluded in our setting.

By the merged pheromone value, the policy generation function is updated yielding a new probability distribution over the arcs: for each $x \in X$ and $a \in A(x)$,

$$\rho_k(x, a) = \frac{\tau_k(x, a)}{\sum_{a' \in A(x)} \tau_k(x, a')}.$$

For improving the performance of the algorithm, *visibility* values $\eta_k(x, a)$ can be introduced into the probability computation: Arcs that seem promising in some *a priori* consideration can be favored. If this option is chosen, one has to replace the last formula above by

$$\rho_k(x, a) = \frac{\tau_k(x, a) \cdot \eta_k(x, a)}{\sum_{a' \in A(x)} \tau_k(x, a') \cdot \eta_k(x, a')},$$

where $\eta_k(x, a) > 0$ is a pre-evaluation value of action a in state x . The simplest way to define $\eta_k(x, a)$ is

$$\eta_k(x, a) = \frac{1}{[C(x, a) + c]^\alpha}$$

for each k , with real numbers $c > 0$, $\alpha > 0$. It is advisable to choose α rather small (say: $\alpha = 0.1$ to 0.5) in order not to bias the learning process governed by the pheromone values $\tau_k(x, a)$ too much towards the “greedy” pre-evaluations $\eta_k(x, a)$. Note that the values $\eta_k(x, a)$ as defined above are always positive.

D. Elite policy generation

The function φ determines how the objective function value of a walk of an ant (i.e., the average value of following a policy) is rewarded into its pheromone. The convergence of Gutjahr’s ant system is based on *elitist strategy* [12] [13] similar to the “*elitism*” concepts introduced by De Jong [9] for GA. In this elitist strategy, the walks that are dominated by another walk that was already traversed do not get pheromone increments anymore. Only the best walks are rewarded.

From the update rule of the pheromone in Section III-A, in order to guarantee the convergence of the algorithm, one has to choose φ such that φ satisfies that for $\pi \in \Lambda_k$, $k = 1, 2, \dots$,

$$\begin{aligned} \varphi(J_\delta^\pi) &> 0 \text{ for } J_\delta^\pi \leq J_\delta^{\pi_{k-1}^e} \\ \varphi(J_\delta^\pi) &= 0 \text{ otherwise.} \end{aligned}$$

We can see that only walks that are at least as good as the performance of the policy π_{k-1}^e get a positive increment for the pheromone. If π_{k-1}^e is equal to the policy that achieves

$$\min_{\pi \in \cup_{i=1}^{k-1} \Lambda_i} J_\delta^\pi,$$

then we have the elitist strategy used in Gutjahr’s ant system. In particular, if a cost function C is positive, we can set $\varphi(J_\delta^\pi) = 1/J_\delta^\pi$ if $J_\delta^\pi \leq J_\delta^{\pi_{k-1}^e}$ and 0 otherwise.

It turns out that we can extend the elitist strategy further. The recent work of policy switching by Chang et al. [6] provides a formal way of combining a finite number of policies in a set to generate an improved policy of all the policies in the set, which we call *elite* policy for the set. It is in general difficult to have a formal rule like policy switching for plain combinatorial optimization problems. Via policy switching, we can change the controlling condition of the elitist strategy so that a walk generated at k th iteration gets rewarded in the pheromone update only if the objective function value of the walk is no worse than the performance of the elite policy of the set $\cup_{i=1}^{k-1} \Lambda_i$. This would make a policy or a walk get rewarded only if the quality of the policy is no worse than any policy generated up to now by ants. By recursively defining the elite policy of k th iteration over the set $\Lambda_k \cup \{\pi_{k-1}^e\}$, the controlling condition of the pheromone update can also include all of the elite policies generated up to now, so that any policy generated by ants so far is included.

Given a nonempty subset Δ of Π , we define a policy $\bar{\pi}$ generated by policy switching with respect to Δ as

$$\bar{\pi}(x) \in \{\arg \min_{\pi \in \Delta} (V^\pi(x))(x)\}, x \in X. \quad (4)$$

We have the following result (see, [7] for a proof):

Theorem 2: Consider a nonempty subset Δ of Π and the policy $\bar{\pi}$ generated by policy switching with respect to Δ given in Equation (4). Then, for all $x \in X$,

$$V^{\bar{\pi}}(x) \leq \min_{\pi \in \Delta} V^\pi(x).$$

The above theorem immediately implies the following result:

Corollary 1: Consider a nonempty subset Δ of Π and the policy $\bar{\pi}$ generated by policy switching with respect to Δ given in Equation (4). Then, for any initial state distribution δ ,

$$J_\delta^{\bar{\pi}} \leq \min_{\pi \in \Delta} J_\delta^\pi.$$

We generate a policy π_k^e , the elite policy with respect to the set $\Lambda_k \cup \{\pi_{k-1}^e\}$, which improves any policy in Λ_k and π_{k-1}^e via policy switching. Note that this is different from the elitist concept of De Jong [9], where the elitist is the best policy in a “population” of $\Lambda_k \cup \{\pi_{k-1}^e\}$. That is, the elitist is a policy that achieves $\min_{\pi \in \Lambda_k \cup \{\pi_{k-1}^e\}} J_\delta^\pi$. Therefore, this immediately implies that the following monotonicity holds:

Corollary 2: For any δ and for all $k \geq 0$,

$$J_\delta^{\pi_k^e} \leq J_\delta^{\pi_{k-1}^e}.$$

The policy switching step is one of the key parts in ANT-PI to *speed up* the convergence of ANT-PI. It can be seen [7] that by one application of policy switching, we eliminate at least $|\Psi|$ number of non-optimal policies but at most $|X|$ number of non-optimal policies in the search process.

Note that policy switching directly manipulates policies to generate an improved policy relative to all policies it was applied to, eliminating the operation of minimization over the entire action space, which is the main computational advantage that replaces the policy improvement step in the original PI.

E. Convergence

Even if the average value of following the two consecutive elite policies are identical, this does not necessarily mean that the elite policy is an optimal policy as in PI. Still, we have a probabilistic convergence guarantee for ANT-PI.

In [16], the following convergence result has been proved. We cite it here within the context of the special situation of the present paper.

Theorem 3: Assume that the following conditions are satisfied:

- (1) There is only one optimal walk through the construction graph.
- (2) Along the optimal walk, the visibility values satisfy $\eta_k(x, a) > 0$.
- (3) Let $J^*(k)$ be the lowest cost value observed in the iterations $1, \dots, k-1$, i.e., the lowest objective function value $J_\delta^{\pi_k}$ corresponding to a walk of an ant in these $k-1$ iterations. (In the case of iteration $k=1$, let $J^*(k) = \infty$.) Then the reward function $\varphi = \varphi_k$ satisfies
 - (i) $\varphi_k(J_\delta^{\pi_k}) > 0$ for $J_\delta^{\pi_k} \leq J^*(k)$,
 - (ii) $\varphi_k(J_\delta^{\pi_k}) = 0$ for $J_\delta^{\pi_k} > J^*(k)$
 for each k .

Moreover, let P_k denote the probability that a fixed ant traverses the optimal walk in iteration k . Then the following two assertions are valid:

- 1) For each $\epsilon > 0$ and for fixed parameter μ , it can be achieved by the choice of a sufficiently large number n of ants that $P_k \geq 1 - \epsilon$ holds for all $k \geq k_0$ (with an integer k_0 depending on ϵ).
- 2) For each $\epsilon > 0$ and for fixed parameter n , it can be achieved by the choice of an evaporation factor μ sufficiently close to zero that $P_k \geq 1 - \epsilon$ holds for all $k \geq k_0$ (with an integer k_0 depending on ϵ).

In [16], this result is proved with the help of four lemmas. Here, we need a slightly stronger version of this theorem (see, [7] for a proof):

Theorem 4: The assertion of Theorem 3 is also valid if condition (i) is replaced by the weaker condition

- (i') $\varphi_k(J_\delta^*) > 0$ for the cost value J_δ^* of the optimal solution.

We are now able to use this general result to obtain a convergence result on ANT-PI (see, [7] for a proof):

Theorem 5: For the algorithm ANT-PI, the following is true: Assume that there is only one optimal policy and that all chosen visibility values are strictly positive, and let P_k denote the probability that ANT-PI produces the optimal policy in iteration k . Then the following two assertions are valid:

- 1) For each $\epsilon > 0$ and for fixed parameter μ , it can be achieved by the choice of a sufficiently large number n of ants that $P_k \geq 1 - \epsilon$ holds for all $k \geq k_0$ (with an integer k_0 depending on ϵ).
- 2) For each $\epsilon > 0$ and for fixed parameter n , it can be achieved by the choice of an evaporation factor μ sufficiently close to zero that $P_k \geq 1 - \epsilon$ holds for all $k \geq k_0$ (with an integer k_0 depending on ϵ).

Note that the condition of the uniqueness of the optimal policy can always be achieved by a very slight change of the cost function that does not make any suboptimal policy appear as optimal: Let $\chi(\pi)$ denote the index of policy π in the chosen fixed order. Add $\epsilon \cdot \chi(\pi)$ to the cost value of each policy π , where ϵ is small enough. Then uniqueness of the optimal solution is obtained.

IV. COMPLEXITY ISSUE

It is known that PI converges faster to the optimal value than VI if both algorithms begin with the same value [21] and that VI is usually outperformed by PI in practical applications. For this reason, we discuss here only complexity issues between ANT-PI and PI.

ANT-PI replaces the single policy evaluation step of PI with multiple policy evaluations, where each policy is generated randomly by an ant, and replaces the single policy improvement step of PI with a combination of a standard ant system framework with policy switching. Doing the single policy evaluation requires solving a system of $|X|$ equations, which takes $O(|X|^3)$ time by Gaussian elimination and $O(|X|^{2.8})$ by the fastest known algorithms. ANT-PI does not add any overhead for this step compared with PI because multiple policy evaluations can be done in parallel independently. At the worst case, it increases the time complexity of policy evaluation by the factor of the number of ants.

As we mentioned before, the policy improvement step takes $O(|X|^2|A|)$ time. If $|A|$ is large, e.g., $O(|X|^2)$, policy improvement takes too much time. If we use an approximation to alleviate this problem, that is, replacing minimization by, e.g., a randomized search, we cannot guarantee the convergence to an optimal policy. On the other hand, policy switching takes $O(n|X|)$ time by directly manipulating policies, where n is the number of ants. Therefore, the main advantage of ANT-PI over PI is for the problems whose action space is large (assuming that the pheromone update is done relatively fast).

The other factor for complexity is the number of iterations required. To the authors' best knowledge, there is no known formal result on this issue for PI except that a

finite number of iterations for the convergence is sufficient (under our finite state and action spaces model). PI may need a very large number of iterations, whereas by the ant system, a good approximative solution may already be reached after a comparatively smaller number of iterations. Note also that the sequence of generated elite policies has a monotonicity property as in the sequence of generated policies in PI. Therefore, if both PI and ANT-PI start with a “good” policy and if PI converges to an optimal policy relatively fast, we expect that ANT-PI also converges to an optimal policy relatively fast.

Suppose that the action space A is not that large, making the direct manipulation of policies by policy switching not helpful. Then, we can replace **Elite Policy Generation** step with a method called “parallel rollout” [6], which is a generalization of PI with *multiple policy* improvement. We simply replace Equation (2) in ANT-PI with

$$\pi_k^e(x) \in \arg \min_{a \in A(x)} \left\{ C(x, a) + \gamma \sum_{y \in X} P(y|x, a) \times \min_{\pi \in \Lambda_k \cup \{\pi_{k-1}^e\}} V^\pi(y) \right\}.$$

Then, every property we discussed before still holds and it is guaranteed that the resulting ANT-PI is faster than PI in terms of the number of iterations (with the same initialization for both algorithms). Also in sequential implementations, it increases the time complexity of policy evaluation only by the factor of the number of ants.

V. CONCLUDING REMARKS

ANT-PI is naturally parallelizable. By doing so, we can improve the runtime. Basically, we partition the policy space Π into subsets of $\{\Pi_i\}$ such that $\bigcup_i \Pi_i = \Pi$ and $\Pi_i \cap \Pi_j = \emptyset$ for all $i \neq j$. Note that there is a corresponding construction graph to each Π_i . We then apply ANT-PI into each Π_i in parallel and then once each part terminates, the best policy π_i^* from each part is taken. We apply then policy switching to the set of best policies $\{\pi_i^*\}$. For a general result regarding parallelization of an algorithm in this respect, see [5].

To obtain a “good” probability bound on the performance of the ant system algorithm in [16], the number of ants and/or evaporation factor needs to be large for applying the algorithm in practice. However, as noted in [16], the analysis is based on a very coarse estimation of the probability that an ant finds the optimal solution so that convergence to the optimum with high probability may be achieved with a moderate number of agents. Following the argument in [16], the purpose of this paper is to demonstrate the capability of the ant-system based approach, in the theoretical perspective, for solving stochastic optimization problems. Certainly, the future works must be implementing and evaluating the proposed algorithms for difficult MDP problems.

REFERENCES

- [1] E. Altman, “Applications of Markov decision processes in communication networks: a survey,” *Markov Decision Processes, Models, Methods, Directions, and Open Problems*, E. Feinberg and A. Shwartz (Eds.) Kluwer, pp. 488–536, 2001.
- [2] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Volumes 1 and 2*. Athena Scientific, 1995.
- [3] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro Dynamic Programming*. Athena Scientific, 1996.
- [4] H. S. Chang, M. Fu, and S. I. Marcus, “An asymptotically efficient algorithm for finite horizon stochastic dynamic programming problems,” in *Proc. of the 42nd IEEE Conf. Decision and Control*, 2003.
- [5] H. S. Chang, H-G. Lee, M. Fu, and S. I. Marcus, “Evolutionary policy iteration for solving Markov decision processes,” *IEEE Trans. on Automatic Control*, revised, 2003.
- [6] H. S. Chang, R. Givan, and E. K. P. Chong, “Parallel rollout for on-line solution of partially observable Markov decision processes,” *Discrete Event Dynamic Systems: Theory and Application*, revised, 2003.
- [7] H. S. Chang, W. J. Gutjahr, J. Yang, and S. Park, “An ant system approach to Markov decision processes,” Tech. Rep. 2003-10, Dept. of Statistics and Decision Support Systems, Univ. of Vienna, Vienna, Austria, 2003.
- [8] P. Cichosz, “Truncating temporal differences: on the efficient implementation of TD(λ) for reinforcement learning,” *Journal of Artificial Intelligence Research*, vol. 2, pp. 287–318, 1995.
- [9] K. A. De Jong, *An Analysis of the Behavior of a Class of Genetic Adaptive Systems*, Ph.D. Thesis, Univ. of Michigan, Ann Arbor, MI, 1975.
- [10] M. Dorigo and G. Di Caro, “The ant colony optimization metaheuristic,” *New Ideas in Optimization*, D. Corne, M. Dorigo (eds.), pp. 11-32, McGraw-Hill, NY, USA, 1999.
- [11] M. Dorigo, G. Di Caro, and T. Stützle, Special Issue on “Ant Algorithms,” *Future Generation Computer Systems*, vol. 16, no. 8, 2000.
- [12] M. Dorigo, V. Maniezzo, A. Colomni, “The Ant System: optimization by a colony of cooperating agents,” *IEEE Trans. Systems Man Cybernet.*, vol. 25, pp. 29–41, 1996.
- [13] L. M. Gambardella and M. Dorigo, “Ant-Q: a reinforcement learning approach to the traveling salesman problem,” *Proc. of the 12th Int. Conf. on Machine Learning*, 1995, pp. 252–260.
- [14] F. Garcia and F. Serre, “Efficient asymptotic approximation in temporal difference learning,” *Proc. of the 14th European Conf. on Artificial Intelligence*, 2000, pp. 296–300.
- [15] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, 1989.
- [16] W. J. Gutjahr, “A graph-based ant system and its convergence,” *Future Generation Computer Systems*, vol. 16, pp. 873–888, 2000.
- [17] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, “Optimization by Simulated Annealing,” *Science*, vol. 220, pp. 45–54, 1983.
- [18] M. Littman, T. Dean, and L. Kaelbling, “On the complexity of solving Markov decision problems,” in *Proc. 11th Annual Conf. on Uncertainty in Artificial Intelligence*, 1995, pp. 394–402.
- [19] Y. Mansour and S. Singh, “On the complexity of policy iteration,” in *Proc. UAI*, 1999, pp. 401–408.
- [20] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, Prentice-Hall, Englewood Cliffs, NJ, 1982.
- [21] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.
- [22] R. Sutton and A. Barto, *Reinforcement Learning*. MIT Press, 2000.