

A LEARNING ALGORITHM FOR MARKOV DECISION PROCESSES WITH ADAPTIVE STATE AGGREGATION ¹

J.S. Baras
Institute for System Research
University of Maryland
College Park, MD 2074, USA
em ail: baras@is.um.edu.

V.S. Borkar
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai 400005, INDIA
em ail: borkar@tifr.res.in.

Abstract

We propose a simulation-based algorithm for learning good policies for a Markov decision process with unknown transition law, with aggregated states. The state aggregation itself can be achieved on a slower time scale by an auxiliary learning algorithm. Rigorous justifications are provided for both algorithms.
KEY WORDS: reinforcement learning, Markov decision processes, learning vector quantization, stochastic approximation, dynamic programming.

1 INTRODUCTION

In classical stochastic control, it is assumed that the state process or an associated observation process is exactly known to the controller. In many realistic situations however, this is not quite true. This is because this information needs to be transmitted across a communication channel with rate constraints, calling for explicit data compression. At one level, one can view compressed observations as further 'partial observations'. But there is one important difference: The data compression scheme itself is a decision variable and should be optimized. Many recent works have addressed these issues [8], [10], [12], [13], [27].

Our aim here is to address this issue in the context of yet another important recent development, viz, reinforcement learning algorithms for learning near-optimal policies (See [3], [22] and the references therein.) These are computationally intensive simulation-based algorithms. They are particularly appealing for large complex systems for which exact analytical modeling and accurate model estimation are unfeasible or expensive,

but simulating typical system transitions is relatively easy (e.g., based on local update rules for interconnected systems). Even so, in their theoretically exact form, these algorithms suffer from the curse of dimensionality, in fact much more so than their classical antecedents like value and policy iteration. Hence there is a need to interface them with a suitable approximation architecture.

The two standard paradigms for approximation are state aggregation and the use of parameterized families of functions (with a low dimensional parameter space) for approximating the value function [3], [24], [25]. The latter, though it takes semantically as an approximation scheme, is not much use from a data compression perspective, whereas state aggregation is natural. This motivates our analysis of a learning algorithm with state aggregation.

To be specific, we consider an 'actor-critic' type learning algorithm. These were introduced in [1] and studied extensively in [14]. We interface this algorithm with state aggregation as proposed in [24], except that this is an 'on line' learning algorithm based on a single simulation run, eschewing in particular the 'independent sampling' over each aggregated state as in [24] which simplifies the analysis of [24] somewhat. Our analysis rigorously justifies aggregating states in terms of the value function. (This was also suggested in [10].) We then propose 'learning' schemes to achieve this aggregation. Taken together, this provides a combined scheme for data compression and optimization.

The paper is organized as follows. The next section formulates our control problem as a Markov decision process on a large but finite state space. It then considers a fixed aggregation of states and formulates the basic algorithm. Section 3 states the main theoretical results in this context, notably the key error estimates. These estimates justify state aggregation on the basis of clustering in the range of the value function. This motivates the clustering algorithm proposed in section 4.

¹This work was initiated when both authors were visiting the Lab. for Information and Decision Systems, M.I.T. Work supported by Grant No. III 5(12)/96-ET from the Department of Science and Technology, Government of India, and U.S. Army Grant PA A L03-92-G-0115, Center for Intelligent Control Systems, M.I.T.

The proofs are quite technical, but a large part of them is standard and is adapted from analogous arguments in [5], [6], [14]. Hence in the interest of brevity we refer the reader to these references for detail, focusing here only on the few points of departure from these references. A complete version of the paper with detailed proofs will appear elsewhere.

2 THE LEARNING PROBLEM

We consider a controlled Markov chain $X(n); n \geq 0$, on a large finite state space $S = \{1, 2, \dots, s\}$, controlled by a control process $Z(n); n \geq 0$, taking values in a finite action space $A = \{a_1, \dots, a_g\}$. (A could be a discretization of a compact metric space.) The transition probability is given by $p: S \times S \times A \rightarrow [0, 1]$. That is

$$P(X(n+1) = i | X(n) = m; Z(n) = a) = p(i; j; a);$$

$fZ(n)g$ is said to be a stationary policy if $Z(n) = v(X(n))$ for some $v: S \rightarrow A$, and a stationary randomized policy if the conditional law of $Z(n)$ given $X(n); Z(n-1); \dots; Z(0)$, is $v(X(n))$ for a $v: S \rightarrow P(A)$. (Here and later, $P(\cdot)$ stands for the space of probability measures on \cdot with topology of weak convergence.) By abuse of terminology, we may refer to the map $v(\cdot)$ (resp., $v(\cdot)$) itself as a stationary (resp. stationary randomized) policy.

Our aim is to minimize over all admissible $fZ(n)g$ the cost

$$E \sum_{m=0}^{\infty} \gamma^m g(X(m); Z(m))$$

for a prescribed running cost $g: S \times R \rightarrow R$ and discount factor $\gamma \in (0, 1)$. The value function $V^* : S \rightarrow R$ associated with this problem is defined as

$$V^*(i) = \min_{fZ(n)g} E \sum_{m=0}^{\infty} \gamma^m g(X(m); Z(m)) | X(0) = i; i \in S;$$

and is the unique solution to the dynamic programming equations

$$V^*(i) = \min_a \left[g(i; a) + \sum_j p(i; j; a) V^*(j) \right], \quad i \in S; \quad (2.1)$$

In fact, $v^* : S \rightarrow P(A)$ is an optimal stationary randomized policy if and only if for each i , $\text{support}(v^*(i)) \subseteq A$ argmin (r.h.s. of (2.1)). A stationary optimal policy always exists obtainable by performing the minimization on the r.h.s. of (2.1). For these and other standard facts concerning Markov decision processes see

[2, 1]. We assume that the chain is irreducible under every stationary randomized policy.

We wish to learn near-optimal policies for an aggregated state space based on a simulation run. The state aggregation is achieved by a partition of the state space into disjoint subsets

$$S = S_1 \cup S_2 \cup \dots \cup S_m; S_i \cap S_j = \emptyset \text{ for } i \neq j;$$

where typically $m \ll s$. We view the S_i 's as aggregated meta-states. Correspondingly, define $p(i; l; a)$, $i \in S, 1 \leq l \leq m, a \in A$, by

$$p(i; l; a) = \sum_{j \in S_l} p(i; j; a);$$

Let $S = \{S_1, S_2, \dots, S_m\}$ and $S^*(n); n \geq 0$, the S -valued process given by $S^*(n) = S_j$ if $X(n) \in S_j; n \geq 0$. Finally, let $f a(n)g; f b(n)g$ be prescribed sequences in $(0, 1)$ such that

$$\begin{aligned} \sum_{n=0}^{\infty} a(n) &= \sum_{n=0}^{\infty} b(n) = 1; \\ \sum_{n=0}^{\infty} a(n)^2 &< 1; \quad \sum_{n=0}^{\infty} b(n)^2 < 1; \\ a(n+1) &= a(n) + \delta; \\ b(n) &= a(n) + \epsilon; \end{aligned}$$

The proposed learning algorithm is as follows. Let $fX(n)g$ be a simulated trajectory of the chain with control sequence $fZ(n)g$ seen to be specified. Let $\tau_0 > 0$ and $j: R \rightarrow [j_0, \tau_0]$ the projection map. The algorithm is a pair of coupled iterations on two different time scales which are simulated by two different choices of step sizes δ, ϵ . The 'fast' component is as follows. For $1 \leq i \leq m$ and $n \geq 0$,

$$\begin{aligned} W_i(n+1) &= W_i(n) + a(n) [f S^*(n) = i g [g(X(n); Z(n)) + \gamma W_{S^*(n+1)}(n) - W_i(n)]]; \quad (2.2) \end{aligned}$$

Here $Z(n) = a$ with probability $\bar{A}_{-i}(n)(S^*(n); a); a \in A$, where

$$\bar{A}_{-i}(i; a) \stackrel{\epsilon}{=} e^{-(i; a)} = \sum_b \bar{A}_{-i}(i; b) e^{-(i; b)}; \quad 1 \leq i \leq m; a \in A;$$

and $\bar{v}^*(n) \stackrel{\delta}{=} [v^*(S^*(n))] \stackrel{\delta}{=} [v^*(i)]$ is given recursively by the 'slow' iteration.

$$\bar{v}^*(i)(n+1) = \sum_j [v^*(j)(n) + b(n) [f S^*(n) = S_j; Z(n) = a g [W_j(n) + g(X(n); a) - \gamma W_{S^*(n+1)}(n)]]]; \quad (2.3)$$

To get an intuitive feel for this scheme, consider $s = m; S_i = \{i\}$ (i.e., no aggregation), whence it reduces to the 'algorithm 3' of [14]. As per the rationale

for two time scale stochastic approximation developed in [4], the fast iteration (2.2) views the slow iteration, i.e. $f^-(n)$, as quasi-static. Thus (2.2) can be approximated by the corresponding iteration with fixed $f^-(n)$ (say). But then it is the stochastic approximation version of fixed policy value iteration, the fixed policy in question being the stationary randomized policy $\pi^-(\Phi)$ given by $\pi^-(i) = \frac{1}{2} \pi^-(i; \Phi; i \in S)$. Thus (2.2) iteratively tracks the fixed policy value function $V^-(\Phi)$ for the policy $\pi^-(\Phi)$, given by

$$V^-(i) = E \sum_{m=0}^{\infty} \alpha^m k(X(m); \zeta(m)) \mid X(0) = i; i \in S;$$

where the expectation is w.r.t. $\pi^-(\cdot)$. Since $\alpha \gg \alpha^-(n)$ varies slowly, it can then be expected to track $V^-(n)$. Plugging this into (2.3) in place of $W(n)$, (2.3) is seen to be a search scheme for minimizing

$$V^-(i; \Phi) + \sum_j \alpha^j p(i; j; \Phi) V^-(j); i \in S;$$

thereby simulating the minimization step of policy iteration. (See [14] for further discussion.) For m as here, our algorithm is then nothing but the algorithm 3 of [14] suitably interfaced with state aggregation.

3 CONVERGENCE ANALYSIS

In this section, we state the main theoretical results concerning the algorithm (2.2)-(2.3). As is always the case with stochastic approximation algorithms their convergence is contingent on their stability, i.e., on their iterates remaining bounded a.s. This is no problem for (2.3) where the projection map ensures the boundedness for free. For (2.2), we need

Lemma 3.1 The iterates of (2.2) remain bounded a.s.

The proof of this lemma is sketched in the appendix. Let $\pi^-(\Phi) \in \mathcal{P}(S)$ denote the unique stationary distribution under the stationary randomized policy $\pi^-(\Phi)$ defined by

$$\pi^-(i) = \frac{1}{2} \pi^-(j; \Phi) \in \mathcal{P}(A) \text{ for } i \in S_j;$$

Define

$$\tilde{A}^-(j) = \sum_{i \in S_i} \pi^-(i) \tilde{A}^-(i; j); j \in S; 1 \leq i \leq m;$$

Lemma 3.2 If we consider (2.2) with the following change: $f^-(n) = \tilde{f}$, a constant in $[i_0^-; \bar{i}_0^-] \times \mathbb{R}^n$, then $W^-(n) \rightarrow W^-$ a.s., where W^- is the unique solution to

$$W^-(i) = \sum_{j \in S_i} \tilde{A}^-(j) \frac{1}{2} \tilde{f}(i; a) g(j; a) + \sum_l \pi^-(j; l; a) W^-(l); 1 \leq i \leq m;$$

The complete proof of this lemma will appear elsewhere. Here we confine ourselves to some important observations concerning (3.1). (3.1) is of the form

$$W^- = G^-(W^-) \quad (3.4)$$

for an appropriately defined $G^-: \mathbb{R}^m \rightarrow \mathbb{R}^m$ satisfying

$$\|G^-(x) - G^-(y)\|_1 \leq \rho \|x - y\|_1; x, y \in \mathbb{R}^m;$$

Here $\|x\|_1 = \sum_i |x_i|$ is the 'max-norm' on \mathbb{R}^m . That (3.2) (equivalently, (3.1)) has a unique solution then follows from the contraction mapping theorem.

Returning to the original setup, we have:

Corollary 3.1 $\|W^-(n) - W^-(\infty)\|_1 \rightarrow 0$ a.s.

Proof This follows from the above lemma and the two time scale argument of [4]. (See, e.g., Lemma 2.3 of [4].)

The two time scale argument of [4] then justifies analysing (2.3) with $W^-(n)$ replacing $W(n)$ on the r.h.s. Let $W^{\pi^-(\Phi)}$ denote the unique solution to

$$W^{\pi^-(\Phi)}(i) = \sum_{j \in S_i} \tilde{A}^-(j) k(j; a) + \sum_l \pi^-(j; l; a) W^{\pi^-(\Phi)}(l); 1 \leq i \leq m;$$

where the existence and uniqueness is ensured as before by the contraction mapping theorem. Then we have

Lemma 3.3 $\|W^-(n) - W^{\pi^-(\Phi)}\|_1 \leq 2 \sum_{i \in S} \pi^-(i) \alpha^i$ a.s., where $\alpha^i \rightarrow 0$ as $i \rightarrow \infty$.

$$\|W^-(i) - W^{\pi^-(\Phi)}(i)\|_1 \leq 2 \sum_{j \in S} \pi^-(j) \alpha^j; \quad (3.5)$$

where $\alpha^i \rightarrow 0$ as $i \rightarrow \infty$.

This follows along the lines of section 5.4 of [14].

Let $\hat{V}^-(j) = W^{\pi^-(\Phi)}(j)$, when $j \in S_i; 1 \leq i \leq m; j \in S$. Let $e = [e_1; \dots; e_m]$ with

$$e_i = \max_{j \in S_i} \hat{V}^-(j) - V^{\pi^-(\Phi)}(j); \quad (3.6)$$

Then we have:

Lemma 3.4 $\|e\|_1 \leq \sum_{i \in S} \pi^-(i) e_i = (1 \mid e)$.

Proof This follows exactly as in Theorem 1 of [24].

Finally, define $V^-(j) = W^-(j)$ when $j \in S_i; 1 \leq i \leq m; j \in S$. The following is then immediate:

Corollary 3.2 $\|V^-(j) - V^{\pi^-(\Phi)}(j)\|_1 \leq (1 \mid e) + 2 \sum_{i \in S} \pi^-(i) \alpha^i$.

The second term can be made arbitrarily small by making the parameter α sufficiently large (though one expects a trade-off with the convergence rate of (2.2)-(2.3)). The first term involves $\sum_{i \in S} \pi^-(i) e_i$, which depends

only on the partition $\{S_i\}$. Given the definition (3.5) of e , it then makes sense to aggregate states based on the value of the value function. The next section proposes a scheme for achieving this

4 CLUSTERING ALGORITHMS

Having concluded that one should aggregate by clustering (or 'vector quantization') in the range of the value function, one important problem remains. The value function is not known. Running a learning algorithm to estimate the value function for sake of state aggregation defeats the purpose, as it brings back into picture the very 'curse of dimensionality' that (2.2)-(2.3) were supposed to work around. Thus a suitable approximation to the value function is called for. Furthermore, we want to come up with a good state aggregation scheme concurrently with (2.2)-(2.3) based on a single simulation run. Hence we are not in a position to first estimate the value function and then aggregate on the basis of it using a batch-mode algorithm like Lloyd's [17]. Once again, two time scale stochastic approximation comes to the rescue: One can run a clustering algorithm on estimated values concurrently, albeit on a slower time scale so as to achieve the desired effect. Implicit in this is the requirement that the algorithm be sequential and based on 'on line' estimates of the value function as they arrive. Thus one has to resort to one of the 'Learning Vector Quantization' (LVQ) schemes. We briefly review them below.

Recall that in vector quantization, one partitions the space into finitely many regions, the 'Voronoi regions', each identified with a point called the 'centroid' therein, as follows. The i -th Voronoi region is the set of all points that are closer to the i -th centroid w.r.t. a suitable metric (usually Euclidean) than any other, any tie being resolved as per some fixed convention. Thus specifying the centroids specifies the Voronoi regions. The connection with clustering is that all points which fall in a single Voronoi region are deemed to form a single cluster. In order that the partition be optimal w.r.t. the mean square error, there is a further requirement. The 'centroids' must in fact be the centroids of the Voronoi regions associated with them w.r.t. the underlying sampling distribution. This however, is a necessary condition and not sufficient, as there can be local minima of the error function that meet this requirement.

The simplest LVQ algorithm is the 'winner take all' competitive learning algorithm described as follows ([11], Ch.9): Let $Q_i(n)$ denote the position of the i -th centroid at time n . Given a new observation $x(n)$, one first locates the centroid that is closest to it (the 'winner'), say $q_j(n)$, and then moves it a little towards

$x(n)$ according to the update

$$q_j(n+1) = (1 - c(n))q_j(n) + c(n)x(n)$$

Here, $c(n) \in (0, 1); \sum_{n=0}^{\infty} c(n) = \infty$, is the usual stochastic approximation-type step size scheme satisfying

$$\sum_{n=0}^{\infty} c(n) = \infty; \quad \sum_{n=0}^{\infty} c(n)^2 < \infty$$

This algorithm can converge to a local minimum of the error function that may be far from the optimal. There are schemes which, while not guaranteeing convergence to the global minimum, do generally improve upon the above algorithm. One such scheme is the Kohonen LVQ algorithm ([11], Ch.9) wherein one updates not only $q_j(n)$, but also the centroids that are its neighbours as per some prescribed neighbourhood scheme. Various 'weighting' schemes have also been proposed to modulate the extent to which different centroids should be moved [18], [20]. A general formalism for such algorithm and its convergence analysis have been presented in [15]. We take this viewpoint here.

Let $q(n) = [q_1(n); \dots; q_m(n)]$ denote the positions of the m centroids at the n -th iteration. The general algorithm is

$$q_i(n+1) = q_i(n) + c(n)I_i(x(n); q(n))(x(n) - q_i(n)); \quad 1 \leq i \leq m$$

Here $I_i(\cdot; \cdot; \cdot)$, $1 \leq i \leq m$, is the 'winner activation function'. For the winner-take all LVQ, $I_i(x; y) = 1$ if $\|x - y\| \leq \|x - y_j\|, \forall j; j \neq i$; 0 otherwise. It can be correspondingly modified for the other LVQ algorithms. The analysis of [15] shows that

$$q(n) \rightarrow q^* = [q_1^*; \dots; q_m^*]$$

with probability one, where each q_i^* is indeed the centroid of the associated Voronoi region w.r.t. the underlying sampling distribution.

Based on this we propose below a scheme for adaptive state aggregation to go with (2.2)-(2.3). Complete theoretical analysis is not presented as it would be quite lengthy, while much of it is standard fare in the analysis of stochastic approximation and related algorithms. This is not to say that there is nothing novel: Once again, it is the points of departure that we shall highlight.

Our scheme is based on the algorithm of Tsitsiklis and Van Roy [25] for learning approximate value function for a fixed policy, based on linear parametrizations. Thus we seek an approximation of the form

$$V_0(i) = \sum_{k=1}^K Q(i; k)h_k; \quad i \in S; \quad (47)$$

where $Q = [Q(i; k)]$ is a prescribed $S \times S$ matrix and $h = [h_1; \dots; h_d]^T$ is the variable parameter vector which we seek to optimize. One chooses $\epsilon > 0$ such that the parameter search is in a space of dimension much lower than $|S| = S$. As before, let $\pi = \frac{1}{2}(\pi^+ + \pi^-)$ be a prescribed stationary randomized policy, and V_π the corresponding π -policy value function. The algorithm of [25] for learning an approximation of V_π of form (42) is (for prescribed $\epsilon \in [0, 1]$):

$$h(n+1) = h(n) + c(n)(g(X(n); Z(n)) + \sum_k Q(X(n+1); k)h_k(n) - \sum_i Q(X(n); k)h_k(n))q(n)$$

$$q(n+1) = \sum_k q(n) + Q(X(n+1); \Phi)$$

where $\sum_k c(n) = 1$; $\sum_k c(n)^2 < 1$.

It is proved in [25] that $h(n) \rightarrow h^*$ with probability one, where h^* is the unique solution to the equations

$$T(\cdot)Qh^* = Qh^*$$

where: Q is the projection operator from R^S onto its d -dimensional subspace spanned by vectors of the form (42), and the operator $T(\cdot) : R^S \rightarrow R^S$ is defined for $0 < \epsilon < 1$ by

$$(T(\cdot)J)(i) = (1 - \epsilon) \sum_{m=0}^{\infty} \epsilon^m E \sum_{n=0}^{\infty} \epsilon^n g(X(n); Z(n)) + \sum_{m=0}^{\infty} \epsilon^{m+1} J(X(m+1)) = X(0) = i$$

and for $\epsilon = 1$ by

$$(T(\cdot)J)(i) = V_\pi(i)$$

$i \in S$. The control sequence $\{Z(n)\}_n$ is realized as per (42).

The aggregation scheme we propose appends (43)-(44) to (2.2)-(2.3), with one difference: $Z(n)$ is now generated with conditional law $\pi_{-i}(\cdot | X(n))$, conditional on $X(n); Z(n-1); m \leq n$. In view of the results of [15] recalled above and the two time scale analysis of [4], it is not hard to show that

$$h(n) \rightarrow h^* \text{ a.s.}$$

Viewing

$$V_\pi^\epsilon = Qh^*$$

as an approximation to V_π , one then aggregates states based on it by algorithm (41), with

$$x(n) = \sum_k Q(X(n); k)h_k^*$$

Of course, the aggregation needs to be done on a much slower time scale than the rest of the algorithm. This is ensured by making $c(n) = \epsilon(n) \rightarrow 0$. To make a correspondence with our earlier notation: Let $S^n = \{S_1^n; \dots; S_m^n\}$ be the partition operative at time n . Then S_i^n is the Voronoi region associated with the centroid $y_i(n)$ for $1 \leq i \leq m$. Invoking the two time scale analysis of [4] in conjunction with the results of the preceding section and of [15], we conclude that

$$q(n) \rightarrow q^* \text{ a.s.}$$

where q^* satisfies the centroid condition w.r.t. π^* for π^* as in (3.4).

It may be recalled that both state aggregation and use of linearly parameterized families were originally advocated as approximation schemes to soften the curse of dimensionality. We combine both, which may seem a bit redundant. The reason for doing so is however, that our concern is not just ameliorating the curse, but also explicit data compression for purposes of transmitting information from the plant to the controller across a communication channel with limited capacity. The state aggregation approach is ideally suited for this as it replaces an S -valued state taking S distinct values by a meta-state that takes only $m \ll S$ distinct values thereby reducing per step bit requirement from $\log_2 S$ to $\log_2 m$. This data compression aspect is not as apparent in the linear parameterization approach, which is more approximation-theoretic in its favour. Our use thereof is for doing a secondary optimization, viz, that over possible aggregation schemes. This is done on a much slower time scale and can be rendered computationally inexpensive by choosing a small value of ϵ . In our formulation, this is of secondary importance to the optimization over policies simplicity in (2.2)-(2.3), hence some laxity in convergence speed and accuracy can be tolerated in the interests of lowering computational overheads.

In conclusion, we have presented a simulation based learning algorithm for Markov decision processes based on the 'actor-critic' paradigm, that explicitly takes into account state aggregation motivated by data compression. A parallel treatment is possible for the alternative learning paradigm of 'Q-learning' [23], [26].

References

- [1] BARTO A., SUTTON R., ANDERSON C., Neuron-like elements that can solve difficult learning control problems IEEE Trans on Systems Man and Cybernetics 13 (1983), 835-844.
- [2] BENVENISTE A., METIVIER M., PRIOUR ET P., Adaptive Algorithms and Stochastic Approximation, Springer Verlag, Berlin-Heidelberg, 1990.

- [3] BERTSEKAS D., TSITSIKLIS J., *Neurodynamic Programming*, Athena Scientific, Belmont, MA, 1996.
- [4] BORKAR V., Stochastic approximation with two time scales *System and Control Letters* 29 (1996), 291-294.
- [5] BORKAR V., Distributed computation of fixed points of 1-nonexpansive maps *Proc. of the Indian Academy of Sciences (Mathematical Sciences)* 106 (1996), 289-300.
- [6] BORKAR V., A synchronous stochastic approximation, *SIAM J. Control and Optim.* 36 (1998), 168-188.
- [7] BORKAR V., MEYN S., The O.D.E. method for convergence of stochastic approximation and reinforcement learning algorithms to appear in *SIAM J. Control and Optim.*
- [8] BORKAR V., MITTER S., LQG control with communication constraints in *Communication, Computation, Control and Signal Processing: A Tribute to Thomas Kailath*, A. Paulraj, V. Roychowdhury, C. Schaper (eds), Kluwer Academic Publ., Norwell, MA, 1997, 365-373.
- [9] BORKAR V., SOUMYANATH K., A new analog parallel scheme for fixed point computation, part I: Theory, *IEEE Trans on Circuits and Systems: Theory and Applications* 44 (1997), 351-355.
- [10] BORKAR V., TATIKONDA S., MITTER S., Markov control problem under communication constraints preprint
- [11] BOSE N.K., LIANG P., *Neural Network Fundamentals with Graphs Algorithm and Applications* McGraw-Hill, Inc., New York, 1996.
- [12] DELCHAMPS D., Stabilizing a linear system with quantized state feedback, *IEEE Trans on Automatic Control* 35 (1990), 916-924.
- [13] FENG X., LOPEZ R., Active probing in control system with quantized state measurements a minimum entropy approach, *IEEE Trans Automatic Control* 42 (1997), 216-238.
- [14] KONDA V.R., BORKAR V., Actor-critic like learning algorithms for Markov decision processes to appear in *SIAM J. Control and Optim.*
- [15] KOSMATOPOULOS E., CHRISTODOULOU M., Convergence properties of a class of Learning Vector Quantization algorithms *IEEE Trans on Image Processing* 5 (1996), 361-368.
- [16] KUSHNER H., YIN G., *Stochastic Approximation Algorithms and Applications* Springer Verlag, New York, 1997.
- [17] LLOYD S., Least square quantization in PCM, *IEEE Trans on Info. Theory* 28 (1982), 129-137.
- [18] MARTINEZ T., BERKOVICH S., SCHULTEN K., 'Neural-gas' network for vector quantization and its application to time-series prediction, *IEEE Trans on Neural Networks* 4 (1993), 553-569.
- [19] NEVEU, J., *Discrete Parameter Martingales* North Holland Amsterdam, 1975.
- [20] PAL N., BEZDEK J., TSAO, E., Generalized clustering networks and Kohonen's self-organizing scheme, *IEEE Trans on Neural Networks* 4 (1993), 549-557.
- [21] PUTERMAN M., *Markov Decision Processes* John Wiley, New York, 1994.
- [22] SUTTON R., BARTO A., *Reinforcement Learning*, MIT Press Cambridge, MA, 1998.
- [23] TSITSIKLIS J., A synchronous stochastic approximation and Q-learning, *Machine Learning*, 16 (1994), 185-202.
- [24] TSITSIKLIS J., VAN ROY, B., Feature-based methods for large scale dynamic programming, *Machine Learning* 22 (1996), 185-202.
- [25] TSITSIKLIS J., VAN ROY B., An analysis of temporal difference learning with function approximation, *IEEE Trans on Automatic Control* 42 (1997), 674-690.
- [26] WATKINS, C., DAYAN, P., Q-learning, *Machine Learning* 8 (1992), 279-292.
- [27] WONG W., BROCKETT R., System with infinite communication bandwidth constraints I: Feedback control problems submitted