



Room 14-0551
77 Massachusetts Avenue
Cambridge, MA 02139
Ph: 617.253.5668 Fax: 617.253.1690
Email: docs@mit.edu
<http://libraries.mit.edu/docs>

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.

Begin text of second and succeeding pages here

Start second column second and succeeding pages here

DISTRIBUTED DYNAMIC PROGRAMMING*ED

Dimitri P. Bertsekas
 Laboratory for Information and Decision Systems
 Massachusetts Institute of Technology
 Cambridge, Massachusetts 02139
 ORGANIZATION

Begin text of first page here. **Abstract** first. Please leave 1/2" space between end of Abstract and first line of text.

We consider distributed algorithms for solving dynamic programming problems whereby several processors participate simultaneously in the computation while maintaining coordination by information exchange via communication links. A model of asynchronous distributed computation is developed which requires very weak assumptions on the ordering of computations, the timing of information exchange, the amount of local information needed at each computation node, and the initial conditions for the algorithm. The class of problems considered is very broad and includes shortest path problems, and finite and infinite horizon stochastic optimal control problems. When specialized to a shortest path problem the algorithm reduces to the algorithm originally implemented for routing of messages in the ARPANET.

1. Introduction

Recent advances in microcomputer technology have intensified interest in distributed computation schemes. Aside from modular expandability, other potential advantages of such schemes are a reduction in computation time for solving a given problem due to parallelism of computation, and elimination of the need to communicate problem data available at geographically dispersed data collection points to a computation center. The first advantage is of crucial importance in real time applications where problem solution time can be an implementation bottleneck. The second advantage manifests itself for example in applications involving communication networks where there is a natural decentralization of problem data acquisition.

The structure of dynamic programming naturally lends itself well to distributed computation since it involves calculations that to a great extent can be carried out in parallel. In fact it is trivial to devise simple schemes taking advantage of this structure whereby the calculation involved in each iteration of the standard form of the algorithm is simply shared by several processors. Such schemes require a certain degree of synchronization

*This research was conducted at the M.I.T. Laboratory for Information and Decision Systems with partial support provided by the National Science Foundation Grant No. NSF/ECS 79-19880.

in that all processors must complete their assigned portion of the computation before a new iteration can begin. As a result complex protocols for algorithm initiation and processor synchronization may be necessary, and the speed of computation is limited to that of the slowest processor. These drawbacks motivate distributed algorithms whereby computation is performed asynchronously at various nodes and independently of the progress in other nodes. Their potential advantages are simpler implementation, faster convergence to a solution and, possibly, a reduction in information exchange between computation nodes.

This paper considers an asynchronous distributed algorithm for a broad class of dynamic programming problems. This class is described in Section 2. The distributed computation model is described in Section 3. It is shown in Section 4 that the algorithm converges to the correct solution under very weak assumptions. For some classes of problems convergence in finite time is demonstrated. These include shortest path problems for which the distributed algorithm of this paper turns out to be essentially the same as the routing algorithm originally implemented in the ARPANET in 1969 [1]. To our knowledge there is no published proof of convergence of this algorithm.

2. Problem Formulation

We use an abstract framework of dynamic programming, first introduced in [2], [3] which includes as special cases a number of specific problems of practical interest.

Let S and C be two sets referred to as the state space and the control space respectively. Elements of S and C are referred to as states and controls and are denoted by x and u respectively. For each $x \in S$ we are given a subset $U(x) \subset C$ referred to as the control constraint set at x . Let F be the set of all extended real valued functions $J: S \rightarrow [-\infty, \infty]$ on S . For any two functions $J_1, J_2 \in F$ we use the notation

$$J_1 \leq J_2 \text{ if } J_1(x) \leq J_2(x), \quad \forall x \in S, \quad (1a)$$

$$J_1 = J_2 \text{ if } J_1(x) = J_2(x), \quad \forall x \in S. \quad (1b)$$

Let $H: S \times C \times F \rightarrow [-\infty, \infty]$ be a mapping which is monotone in the sense that for all $x \in S$ and $u \in U(x)$ we have

$$H(x,u,J_1) \leq H(x,u,J_2) \text{ for } \forall J_1, J_2 \in F, \text{ with } J_1 \leq J_2. \quad (2)$$

Given a subset $\bar{F} \subset F$ the problem is to find a function $J^* \in \bar{F}$ such that

$$J^*(x) = \inf_{u \in U(x)} H(x,u,J^*), \quad \forall x \in S. \quad (3)$$

By considering the mapping $T: F \rightarrow F$ defined by

$$T(J)(x) = \inf_{u \in U(x)} H(x,u,J) \quad (4)$$

the problem is alternately stated as one of finding a fixed point of T within \bar{F} , i.e., a function $J^* \in \bar{F}$ such that

$$J^* = T(J^*). \quad (5)$$

We will assume throughout that T has a unique fixed point within \bar{F} .

BEGIN PAPER TWO SPACES

We provide some examples that illustrate the broad scope of the problem formulation just given.

Example 1 (Shortest Path Problems): Let (N,L) be a directed graph where $N = \{1,2,\dots,n\}$ denotes the set of nodes and L denotes the set of links. Let $N(i)$ denote the downstream neighbors of node i , i.e., the set of nodes j for which (i,j) is a link. Assume that each link (i,j) is assigned a positive scalar a_{ij} referred to as its length. Assume also that there is a directed path to node 1 from every other node. Then it is known ([4], p.67) that the shortest path distances d_i^* to node 1 from all other nodes i solve uniquely the equations

$$d_i^* = \min_{j \in N(i)} \{a_{ij} + d_j^*\}, \quad \forall i \neq 1 \quad (6a)$$

$$d_1^* = 0. \quad (6b)$$

If we make the identifications

$$S = C = N, U(x) = N(x), \bar{F} = F, J^*(x) = d_x^* \quad (7)$$

$$H(x,u,J) = \begin{cases} a_{xu} + J(u) & \text{if } x \neq 1 \\ 0 & \text{if } x = 1 \end{cases} \quad (8)$$

we find that the abstract problem (3) reduces to the shortest path problem.

Example 2 (Infinite Horizon Stochastic Optimal Control Problems): Let H be given by

$$H(x,u,J) = E\{g(x,u,w) + \alpha J[f(x,u,w)] | x,u\} \quad (9)$$

where the following are assumed:

(1) The parameter w takes values in a countable set W with given probability distribution $p(dw|x,u)$ depending on x and u , and $E\{\cdot | x,u\}$ denotes expected value with respect to this distribution.

(2) The functions g and f map $S \times C \times W$ into $[-\infty, \infty]$ and S respectively.

(3) The scalar α is positive.

Because the set W is assumed countable the expected value in (9) is well defined for all $J \in F$ in terms of infinite summation provided we use the convention $+\infty \cdot \infty = +\infty$ (see [3], p.31). It is possible to consider a more general probabilistic structure for W (see [3]) at the expense of complicating the presentation but this does not seem worthwhile in view of the computational purposes of the paper.

Begin second column of text here on first page only.

It is shown in [3] that with this definition of H the abstract problem (3) reduces under more specific assumptions to various types of standard stochastic optimal control problems. Thus if g is uniformly bounded above and below by some scalars and $0 < \alpha < 1$ the problem is equivalent to the standard infinite horizon discounted stochastic optimal control problem with bounded cost per stage (see [5] Sections 6.1-6.3). Under these circumstances the mapping T of (4) has a unique fixed point J^* in the class of all bounded real valued functions on S and J^* is the optimal value function of the corresponding stochastic optimal control problem.

If we assume that $0 \leq g(x,u,w)$ or $g(x,u,w) \leq 0$ for all $(x,u,w) \in S \times C \times W$ then we obtain stochastic optimal control problems of the type discussed extensively, for example, in [5], Sections 6.4-6.6, 7.1-7.4, and [3], Chapter 5. If J^* is the optimal value function for such a problem then J^* is the unique fixed point of T over all functions $J \in F$ such that $0 \leq J < J^*$ if $0 < g(x,u,w)$ for all (x,u,w) , or $J^* < J \leq 0$ if $g(x,u,w) \leq 0$ for all (x,u,w) (see [5], p. 256).

Example 3 (Finite Horizon Stochastic Optimal Control Problems): Let $S, C, U(x), W, p(dw|x,u), g$ and f be as in Example 2 and consider the set of equations

$$J_N(x_N) = 0, \quad x_N \in S \quad (10a)$$

$$J_k(x_k) = \inf_{u_k \in U(x_k)} E\{g(x_k, u_k, w_k) + J_{k+1}[f(x_k, u_k, w_k)] | x_k, u_k\}, \quad (10b)$$

$$k = 0, 1, \dots, N-1, x_k \in S,$$

where N is a positive integer. These are the usual dynamic programming equations associated with finite horizon stochastic optimal control problems with zero terminal cost and stationary cost per stage and system function. It is possible to write these equations in the form (5) by defining a new state space consisting of an $(N+1)$ -fold Cartesian product of S with itself, writing $J^* = (J_0, J_1, \dots, J_N)$, and appropriately defining H on the basis of

(10). In fact this is a standard procedure for converting a finite horizon problem to an infinite horizon problem (see [5], p.325). This reformulation

can also be trivially generalized to finite horizon problems involving a nonzero terminal cost and a nonstationary system and cost per stage.

3. A Model for Distributed Dynamic Programming

Our algorithm can be described in terms of a collection of n computation centers referred to as nodes and denoted $1, 2, \dots, n$. The state space S is partitioned into n disjoint sets denoted S_1, \dots, S_n .

Each node i is assigned the responsibility of computing the values of the solution function J^* at all states x in the corresponding set S_i . A node

j is said to be a neighbor of node i if $j \neq i$ and there exist a state $x_i \in S_i$ and two functions $J_1, J_2 \in F$ such that

$$J_1(x) = J_2(x), \quad \forall x \in S_j \quad (11a)$$

$$T(J_1)(x_i) \neq T(J_2)(x_i) \quad (11b)$$

The set of all neighbors of i is denoted $N(i)$. Intuitively j is not a neighbor of i if, for every $J \in F$, the values of J on S_j do not influence the values of $T(J)$ on S_i . As a result, for any $J \in F$, in order for node i to be able to compute $T(J)$ on S_i it is only necessary to know the values of J on sets $S_j, j \in N(i)$, and, possibly, on the set S_i .

At each time instant, node i can be in one of three possible states compute, transmit, or idle. In the compute state node i computes a new estimate of the values of the solution function J^* for all states $x \in S_i$. In the transmit state node i communicates the estimate obtained from the latest compute phase to one or more nodes m for which $i \in N(m)$. In the idle state node i does nothing related to the solution of the problem. It is assumed that a node can receive a transmission from neighbors simultaneously with computing or transmitting, but this is not a real restriction since, if needed, a time period in a separate receive state can be lumped into a time period in the idle state.

We assume that computation and transmission for each node takes place in uninterrupted time intervals $[t_1, t_2]$ with $t_1 < t_2$, but do not exclude the possibility that a node may be simultaneously transmitting to more than one nodes nor do we assume that the transmission intervals to these nodes have the same origin and/or termination. We also make no assumptions on the length, timing and sequencing of computation and transmission intervals other than the following:

Assumption (A): There exists a positive scalar P such that, for every node i , every time interval of length P contains at least one computation interval for i and at least one transmission interval from i to each node m with $i \in N(m)$.

Each node i also has one buffer per neighbor $j \in N(i)$ denoted B_{ij} where it stores the latest trans-

mission from j , as well as a buffer B_{ii} where it stores its own estimate of values of the solution function J^* for all states $x \in S_i$. The contents of each buffer B_{ij} where $j=i$ or $j \in N(i)$ at time t are

denoted J_{ij}^t . Thus J_{ij}^t is, for every t , a function from S_j into $[-\infty, \infty]$ and may be viewed as the estimate by node i of the restriction of the solution function J^* on S_j available at time t . The rules according to which functions J_{ij}^t are updated are as follows:

1) If $[t_1, t_2]$ is a transmission interval for node j to node i with $i \in N(j)$ the contents $J_{jj}^{t_1}$ of the buffer B_{jj} at time t_1 are transmitted and entered in the buffer B_{ij} at time t_2 , i.e.

$$J_{ij}^{t_2} = J_{jj}^{t_1} \quad (12)$$

2) If $[t_1, t_2]$ is a computation interval for node i the contents of buffer B_{ii} at time t_2 are replaced by the restriction of the function $T(J_{ij}^{t_1})$ on S_i where, for all t , J_{ij}^t is defined by

$$J_{ij}^t(x) = \begin{cases} J_{ii}^t(x) & \text{if } x \in S_i \\ J_{ij}^t(x) & \text{if } x \in S_j \text{ and } j \in N(i) \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

In other words we have

$$J_{ii}^{t_2}(x) = T(J_{ij}^{t_1})(x) = \inf_{u \in U(x)} H(x, u, J_{ij}^{t_1}), \quad \forall x \in S_i \quad (14)$$

3) The contents of a buffer B_{ii} can change only at the end of a computation interval for node i . The contents of a buffer $B_{ij}, j \in N(i)$, can change only at the end of a transmission interval from j to i .

Note that by definition of the neighbor set $N(i)$, the value $T(J_{ij}^t)(x)$ for $x \in S_i$ does not depend on the values of J_{ij}^t at states $x \in S_m$ with $m \neq i$, and $m \in N(i)$. We have assigned arbitrarily the default value zero to these states in (13). Our objective is to show that for all $i = 1, \dots, n$

$$\lim_{t \rightarrow \infty} J_{ij}^t(x) = J^*(x), \quad \forall x \in S_j, \quad j = i \text{ or } j \in N(i).$$

It is clear that an assumption such as (A) is necessary in order for such a result to hold. Since iteration (14) is of the dynamic programming type it is also clear that some restrictions must be placed on the mapping H that guarantee convergence of the algorithm under the usual circumstances where the algorithm is carried out in a centralized, synchro-

nous manner (i.e., when there is only one computation node). The following assumption places somewhat indirect restrictions on H but simplifies the convergence analysis.

Assumption (B): There exist two functions \underline{J} and \bar{J} in F such that the set of all functions $J \in F$ with $\underline{J} \leq J \leq \bar{J}$ belongs to F and furthermore

$$J \geq T(\bar{J}), \quad T(\underline{J}) \geq \underline{J} \quad (15)$$

$$\lim_{k \rightarrow \infty} T^k(\bar{J})(x) = J^*(x), \quad \forall x \in S$$

and $\lim_{k \rightarrow \infty} T^k(\underline{J})(x) = J^*(x), \quad \forall x \in S$. Please leave (16) space between end of Abstract and first line of text.

where T^k denotes composition of the mapping T with itself k times.

Note that in view of the monotonicity of H [cf. (2)] and the fact $J^* = T(J^*)$, Assumption (B) implies

$$\bar{J} \geq T(\bar{J}) \geq T^2(\bar{J}) \geq \dots \geq J^* \geq \dots$$

$$\dots \geq T^2(\underline{J}) > T(\underline{J}) > \underline{J}$$

Furthermore if $J \in F$ satisfies $\underline{J} \leq J \leq \bar{J}$ then

$$\lim_{k \rightarrow \infty} T^k(J)(x) = J^*(x) \text{ for all } x \in S.$$

Assumption (B) leaves open the question of how to find suitable functions \underline{J} and \bar{J} . On the other hand for most problems of interest the choices are clear. In particular we have the following:

1) For shortest path problems (Example 1) it is straightforward to verify that the choices

$$\underline{J}(i) = 0, \quad \forall i = 1, \dots, n \quad (17a)$$

$$\bar{J}(i) = \begin{cases} \infty & \text{if } i \neq 1 \\ 0 & \text{if } i = 1 \end{cases} \quad (17b)$$

satisfy Assumption (B).

2) For finite horizon stochastic optimal control problems (Example 3) for which the function g in (10) is uniformly bounded below it can be easily verified that the functions $\bar{J} = (\bar{J}_0, \bar{J}_1, \dots, \bar{J}_N)$ and $\underline{J} = (\underline{J}_0, \underline{J}_1, \dots, \underline{J}_N)$ where for all k and x

$$\bar{J}_k(x) = \infty, \quad \underline{J}_k(x) = -\infty \quad (18)$$

satisfy Assumption (B).

3) For discounted infinite horizon stochastic optimal control problems with bounded cost per stage (Example 2 with $\alpha \in (0, 1)$ and g uniformly bounded above and below) it is easily shown that every pair of functions \bar{J}, \underline{J} of the form

$$\bar{J}(x) = \bar{\beta}, \quad \underline{J}(x) = \underline{\beta}, \quad \forall x \in S$$

where

$$\frac{1}{1-\alpha} (\sup_{x \in S} \inf_{u \in U(x)} E\{g(x, u, w)\}) \leq \bar{\beta} \leq \frac{1}{1-\alpha} (\inf_{x \in S} \sup_{u \in U(x)} E\{g(x, u, w)\})$$

$$-\infty < \underline{\beta} \leq \frac{1}{1-\alpha} \inf_{x \in S} \inf_{u \in U(x)} E\{g(x, u, w)\}$$

satisfy Assumption (B).

4) For infinite horizon stochastic optimal control problems with nonpositive cost per stage (Example 2 with $g \leq 0$) it can be shown that the functions \bar{J}, \underline{J} with

$$\underline{J}(x) = J^*(x), \quad \bar{J}(x) = 0, \quad \forall x \in S$$

satisfy Assumption (B) ([5], pp. 261, 298). If the cost per stage is nonnegative (Example 2 with $g \geq 0$) then, under a mild assumption (which is satisfied in particular if $U(x)$ is a finite set for each x), it can be shown that the choices \bar{J}, \underline{J} with

$$\underline{J}(x) = 0, \quad \bar{J}(x) = J^*(x), \quad \forall x \in S$$

satisfy Assumption (B) ([5], pp. 263, 298). The choice of \underline{J} and \bar{J} can be further sharpened and simplified under more specific assumptions on problem structure but we will not pursue this matter further.

Our convergence result will be shown under the assumption that the contents J_{ij}^0 of the buffers B_{ij} at the initial time $t=0$ satisfy

$$\underline{J}(x) \leq J_{ij}^0(x) \leq \bar{J}(x), \quad \forall x \in S_j. \quad (19)$$

The broad range of initial conditions allowed by (19) eliminates the need to reset the contents of the buffers in an environment where it is necessary to execute periodically the algorithm with slightly different problem data as for example in routing algorithms for communication networks. This is particularly true for cases 1)-3) above where condition (19) implies that the initial buffer contents can be essentially arbitrary.

4. Convergence Analysis

Our main result is the following proposition.

Proposition 1: Let Assumptions (A) and (B) hold and assume that for all $i = 1, \dots, n$

$$\underline{J}(x) \leq J_{ij}^0(x) \leq \bar{J}(x), \quad \forall x \in S_j, \quad j = i \text{ or } j \in N(i). \quad (20)$$

Then for all $i = 1, \dots, n$

$$\lim_{t \rightarrow \infty} J_{ij}^t(x) = J^*(x), \quad \forall x \in S_j, \quad j = i \text{ or } j \in N(i).$$

Proof: Since the contents of a buffer can change only at the end of a computation or transmission interval at a node we can analyze convergence in terms of a discrete time process. We focus attention at the sequence of times $\{t_k\}$ with $0 < t_1 < t_2 < \dots$ where each t_k is the end of a computation interval for one or more nodes.

Let us consider for all $t \geq 0$

$\bar{J}_{ij}^t(x)$ for $x \in S_j$ and $t \in [0, \infty]$; $\bar{J}_{ij}^t(x)$ for $x \in S_j$ and $t \in [0, \infty]$, here
 where for each $x \in S_j$, the value $\bar{J}_{ij}^t(x)$ [$\bar{J}_{ij}^t(x)$] represents the contents of buffer B_{ij} at time t if the algorithm were executed with the same timing and order of computation and transmission intervals but with the initial condition $\bar{J}(x)$ [$\underline{J}(x)$] instead of $J_{ij}^0(x)$ for each buffer B_{ij} and $x \in S_j$. The monotonicity assumption (2) and the definition of the algorithm [cf. (13), (14)] clearly imply that for all

$$\bar{J}_{ij}^t(x) \leq \bar{J}_{ij}^{t_1}(x) \leq \bar{J}_{ij}^{t_2}(x), \quad \forall x \in S_j, \quad t_1 < t_2, \quad i = 1, \dots, n, \quad j = i \text{ or } j \in N(i). \quad (21)$$

It will thus suffice to show that

$$\lim_{t \rightarrow \infty} \bar{J}_{ij}^t(x) = J^*(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i) \quad (22)$$

$$\lim_{t \rightarrow \infty} \bar{J}_{ij}^t(x) = \underline{J}^*(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i) \quad (23)$$

In view of the fact $\bar{J} \geq T(\bar{J})$ we have clearly

$$\bar{J}(x) = \bar{J}_{ij}^0(x) \geq \bar{J}_{ij}^{t_1}(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i) \quad (24)$$

with potential strict inequality only for nodes j for which t_1 was the end of a computation interval. For $t \in [t_1, t_2)$ the content of B_{ij} is either \bar{J}_{ij}^0 or $\bar{J}_{jj}^{t_1}$ so from (24) we must have

$$\bar{J}_{ij}^t(x) \geq \bar{J}_{jj}^{t_1}(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i), \quad t \in [t_1, t_2) \quad (25)$$

$$\bar{J}_{ij}^{t_1}(x) \geq \bar{J}_{ij}^t(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i), \quad t \in [t_1, t_2).$$

The last relation can also be written as

$$\bar{J}_{jm}^{t_1}(x) \geq \bar{J}_{jm}^t(x), \quad \forall x \in S_m, \quad j = 1, \dots, n, \quad m \in N(j), \quad t \in [t_1, t_2). \quad (26)$$

In view of the monotonicity of H it follows from (26) and (14) that

$$\bar{J}_{jj}^{t_1}(x) \geq \bar{J}_{jj}^{t_2}(x), \quad \forall x \in S_j \quad (27)$$

with potential strict inequality only for nodes j for which t_2 was the end of a computation interval. Combining (25) and (27) we obtain

$\bar{J}_{ij}^t(x) \geq \bar{J}_{jj}^t(x)$, $\forall x \in S_j$, $i = 1, \dots, n$, $j \in N(i)$
 with potential strict inequality only for nodes j for which either t_1 or t_2 was the end of a computation interval. The preceding argument can be repeated to show that for all k , $i = 1, \dots, n$, and $j \in N(i)$ we have

$$\bar{J}_{ij}^{t_k}(x) > \bar{J}_{ij}^{t_{k-1}}(x) > \bar{J}_{jj}^{t_k}(x) > \bar{J}_{jj}^{t_{k-1}}(x), \quad \forall x \in S_j, \quad t \in [t_k, t_{k+1}). \quad (28)$$

Let k_1 be the first integer for which $2P < t_{k_1}$ where P is as in Assumption (A). Then each node k_1 must have completed at least one computation phase in the interval $[0, P]$ and at least one transmission phase to all nodes in the interval $[P, 2P]$. It is easily seen that this together with (28), the monotonicity of H , and the definition of the algorithm implies that for all $t \in [t_{k_1}, t_{k_1+1})$

$$T(\bar{J})(x) \geq \bar{J}_{ij}^{t_{k_1}}(x) \geq \bar{J}_{jj}^{t_{k_1}}(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i).$$

This argument can be repeated and shows that if $m(k)$ is the largest integer m such that $2mP \leq t_k$ then for all $t \in [t_k, t_{k+1})$

$$T^{m(k)}(\bar{J})(x) \geq \bar{J}_{ij}^t(x) \geq \bar{J}_{jj}^{t_k}(x), \quad \forall x \in S_j, \quad j \in N(i). \quad (29)$$

Similarly we obtain for all $t \in [t_k, t_{k+1})$

$$\bar{J}_{jj}^{t_k}(x) \geq \bar{J}_{ij}^t(x) \geq T^{m(k)}(\underline{J})(x), \quad \forall x \in S_j, \quad i = 1, \dots, n, \quad j \in N(i). \quad (30)$$

By combining (21), (29), and (30) and using Assumption (B) we obtain (22), (23) and the proof of the proposition is complete. Q.E.D.

Note that (21), (29), and (30) provide useful estimates of rate of convergence. In fact by using these relations it is possible to show that in some special cases convergence is attained in finite time.

Proposition 2: In the cases of Examples 1 and 3 with the choices of \bar{J} , \underline{J} given by (17) and (18) respectively and J_{ij}^0 satisfying (20), there exists a time $\bar{t} > 0$ such that for all $i = 1, \dots, n$ and $t \geq \bar{t}$ there holds

$$\bar{J}_{ij}^t(x) = J^*(x), \quad \forall x \in S_j, \quad j = i \text{ or } j \in N(i).$$

Proof: For Example 3 it is easily seen that there holds

$$T^k(\bar{J})(x) = T^k(\underline{J})(x) = J^*(x), \quad \forall x \in S, \quad k \geq N+1$$

The proof follows from (21), (29) and (30). For example 1 it is easily seen that

$$T^k(J)(i) = J^*(i), \quad i = 2, \dots, n, \quad k \geq k_0 \text{ ON FIRST PAGE HERE, CENTERED}$$

Also for each i , $T^k(J)(i)$ represents the length of a path starting from i with k links; and each link has positive length. Therefore there exists a k such that $T^k(J)(i)$ represents length of a path from i to node 1, for otherwise the paths corresponding to $T^k(J)(i)$ would cycle indefinitely without reaching node 1 and we would have $T^k(J)(i) \rightarrow \infty$. Since $T^k(J)(i) \leq J^*(i)$ and $J^*(i)$ is the shortest distance from i to 1 we obtain

$$T^k(J)(i) = J^*(i), \quad \forall i = 2, \dots, n, \quad k \geq k_0.$$

The result again follows from (21), (29) and (30).
Q.E.D.

It is possible to construct examples showing that in the case of the shortest path problem the number of iterations needed for finite convergence of $T^k(J)$ depends on the link lengths in a manner which makes the overall algorithm nonpolynomial.

In many problems of interest the main objective of the algorithm is to obtain a minimizing control law μ^* , i.e. a function $\mu^*: S \rightarrow C$ with $\mu^*(x) \in U(x)$ for all $x \in S$ such that

$$H[x, \mu^*(x), J^*] = \min_{u \in U(x)} H(x, u, J^*), \quad \forall x \in S. \quad (31)$$

It is thus of interest to investigate the question of whether control laws $\mu^t: S \rightarrow C$ satisfying

$$\mu^t(x) \in U(x), \quad \forall x \in S \quad (32)$$

and

$$H[x, \mu^t(x), J_i^t] = \min_{u \in U(x)} H(x, u, J_i^t), \quad \forall x \in S_i, \quad i=1, \dots, n \quad (33)$$

where J_i^t is given for all t by (13), converge in some sense to a control law μ^* satisfying (31). The following proposition shows that convergence is attained in finite time if the sets $U(x)$ are finite and H has a continuity property which is satisfied for most problems of practical interest. A related convergence result can be shown assuming the sets $U(x)$ are compact (c.f. [5], Prop. 5.11).

Proposition 3: Let the assumptions of Proposition 1 hold. Assume also that for every $x \in S$, $u \in U(x)$ and sequence $\{J^k\} \subset \bar{F}$ for which $\lim_k J^k(x) = J^*(x)$ for all $x \in S$ we have

$$\lim_{k \rightarrow \infty} H(x, u, J^k) = H(x, u, J^*). \quad (34)$$

Then for each state $x \in S$ for which $U(x)$ is a finite set there exists $\bar{t}_x > 0$ such that for all $t \geq \bar{t}_x$ if $\mu^t(x)$ satisfies (32), (33) then

$$H[x, \mu^t(x), J^*] = \min_{u \in U(x)} H(x, u, J^*).$$

Proof: See [6]; second and succeeding pages here

5. Conclusions

The analysis of this paper shows that natural distributed dynamic programming schemes converge to the correct solution under very weak assumptions on the problem structure, and the timing and ordering of computation and internode communication. The restrictions on the initial conditions are also very weak. This means that, for problems that are being solved continuously in real time, it is not necessary to reset the initial conditions and resynchronize the algorithm each time the problem data changes. As a result the potential for tracking slow variations in optimal control laws is improved, and algorithmic implementation is greatly simplified.

References

- [1] J. McQuillan, G. Falk, and I. Richer, "A Review of the Development and Performance of the ARPANET Routing Algorithm", *IEEE Trans. on Communications*, Vol. COM-26, 1978, pp. 1802-1811.
- [2] D. P. Bertsekas, "Monotone Mappings with Application in Dynamic Programming", *SIAM J. Control and Optimization*, Vol. 15, 1977, pp. 438-464.
- [3] D. P. Bertsekas, and S. E. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*, Academic Press, N.Y., 1978.
- [4] E. L. Lawler, *Combinatorial Optimization: Networks and Matroids*, Holt, Rinehart, and Winston, N.Y., 1976.
- [5] D. P. Bertsekas, *Dynamic Programming and Stochastic Control*, Academic Press, N.Y., 1976.
- [6] D. P. Bertsekas, "Distributed Dynamic Programming", Report LIDS-P-1060, Lab. for Information and Decision Systems, M.I.T., Cambridge, MA, 1980 (to appear in *IEEE Trans. on Aut. Control*).