Online Companion for

"Decentralized Adaptive Flow Control of High Speed Connectionless Data Networks"

Operations Research
Vol. 47, No. 6
November-December 1999

Felisa J. Vásquez-Abad
Départment d'informatique et Recherche Opérationnelle
Université de Montréal
Montréal, Québec, Canada

Lorne G. Mason
INRS-Télécommunications
Université du Québec
Verdun, Québec, Canada

Appendix

Appendix A. The Proof of Asymptotic Optimality.

We present the proof of convergence of the algorithms (12) and (14), using the framework and results of the weak convergence method in [8] and [21].

Call τ_n the epoch of the n-th *global* update, regardless of the controller, and τ_n^c the epoch of the n-th local update at c. Denote by θ the vector of all the distribution probabilities π_i^c. Then both algorithms can be expressed as a stochastic approximation:

$$\theta_{n+1}^\epsilon = \theta_n^\epsilon + \epsilon Y_n^\epsilon,$$

if we identify $\epsilon = G$ with the gain parameter, and n with the n-th global update epoch. Clearly, the vector Y_n^ϵ depends on the feedback function $b(\tau_n)$. The process can be imbedded in a Markov Decision Process (MDP) $(\xi_n^\epsilon, \theta_n^\epsilon)$, identifying the state ξ_n^ϵ with the vector of queue lengths, residual service times (if not Markovian) and local information at time τ_n. Following the notation in [8], call $G(\xi,\theta) = E\{Y_n^\epsilon | \xi_n^\epsilon = \xi, \theta_n^\epsilon = \theta\}$. The σ-algebra related to the MDP up to the n-th update will be denoted by $\mathcal{F}_n^\epsilon$. By the Markovian property, $G(\theta_n^\epsilon, \xi_n^\epsilon) = E\{Y_n^\epsilon | \mathcal{F}_n^\epsilon\}$, which is a random variable depending on the distribution of $(\xi_n^\epsilon, \theta_n^\epsilon)$. From the closed network model, it follows that the fixed control process $\xi(\theta)$ Markovian with transition probability $P(dx, x) = P\{\xi_{n+1}(\theta) \in dx | \xi_n(\theta) = x\}$, which is weakly continuous in θ for our model. A closed network with stationary service distributions (independent of θ)) is stable for every possible value of θ. Call $\mu_\theta(dx)$ the invariant measure of the fixed control process, and

$$g(\theta) = \int \mu_\theta(dx) G(x,\theta) = \lim_{m\to\infty} \frac{1}{m} \sum_{n=0}^{m-1} E\{G(\xi_n(\theta),\theta)\}$$

Notice that this latter expectation is w.r.t. the fixed control process. The random variables Y_n^ϵ are uniformly integrable, since they are uniformly bounded by construction of the feedback functions. Furthermore, the sequence $\{(\xi_n^\epsilon, \theta_n^\epsilon), n \geq 0, \epsilon > 0\}$ is tight, that is, for every $\alpha > 0$ there exist a compact set K_α such that $\sup_{\epsilon,n} P\{(\xi_n^\epsilon, \theta_n^\epsilon) \in K_\alpha\} < \alpha$. This follows because θ_n^ϵ are probability vectors, and ξ_n^ϵ are uniformly tight: queue sizes are all bounded by the window size, the residual services are independent of $\epsilon, \theta_n^\epsilon$ and the information vector containing the moving averages is also bounded a.s. by a random variable of maximal delay along feasible paths, which is independent of θ_n^ϵ and n. Tightness is the stochastic analog of compactness. Define the control interpolation process:

$$\vartheta^\epsilon(t) = \theta_n^\epsilon \quad \text{for } t \in [n\epsilon, (n+1)\epsilon) \tag{A.1}$$

From this definition and the stochastic approximation form, for $t = \epsilon n$ we have:

$$\frac{\vartheta^\epsilon(t+\epsilon) - \vartheta^\epsilon(t)}{\epsilon} = Y_n^\epsilon$$

and the conditional expected behaviour of Section 4 is related to:

$$E\left(\frac{\vartheta^\epsilon(t+\epsilon) - \vartheta^\epsilon(t)}{\epsilon} | \mathcal{F}_n^\epsilon\right) = G(\vartheta^\epsilon(t), \xi_n^\epsilon)$$

From Proposition 1 in [21] it follows that every subsequence of $\vartheta^\epsilon(\cdot)$ as $\epsilon \to 0$ has a further weakly convergent subsequence and all weak limits are Lipshitz continuous a.s. All the assumptions of [8] are satisfied, therefore any such limit satisfies the ODE:

$$\frac{d\vartheta(t)}{dt} = g[\vartheta(t)]$$

If the ODE has a unique solution for each initial condition, the limit does not depend on the subsequence and therefore $\vartheta^\epsilon(\cdot)$ converges to $\vartheta(\cdot)$. It is common to assume that $g(\cdot)$ is locally Lipschitz continuous, thus continuous and uniformly bounded on compact sets, which would ensure uniqueness of the solution for each initial condition. If, furthermore, this ODE has asymptotically stable points $\theta^* \in \mathcal{S}$, then the limit points $\lim_{t\to\infty} \vartheta(t)$ of its solutions (which may or may not depend on the initial condition) belong to $\mathcal{S}$ and satisfy

$g(\theta^*) = 0$. In particular, for the schemes presented in the present work, it follows from (11) that, for the fixed control process:

$$F_i(\theta) = \lim_{n \to \infty} \frac{1}{n} \sum_{n=0}^{m-1} E\left\{ E\{F_i^c(\tau_n^c) | \mathcal{F}_n^c(\theta)\}\right\} = 1 - \frac{\arctan H_i[X(\theta)]}{(\pi/2)} \tag{A.2}$$

for all c, where $\mathcal{F}_n^c(\theta)$ is the σ-algebra generated by $\{\xi_n(\theta)\}$ up to time τ_n^c. Let $\vartheta_i^c(t)$ denote the limit process of the component corresponding to the control variable π_i^c. Then it follows that the limit process corresponding to the scheme in (12) satisfies:

$$\frac{d\vartheta_i^c(t)}{dt} = M_d[\vartheta(t)] \left[F_i[\vartheta(t)] - \vartheta_i^c(t) \sum_{o=1}^{N} F_o[\vartheta(t)] \right] \tag{A.3}$$

where $M_d(\theta)$ is the stationary fraction of updates performed at controller d for the fixed control process (see the proof of Theorem 3 in [21] for the details on the time scaling argument in decentralized schemes). Analogously, for the scheme in (14), the limit processes satisfy:

$$\frac{d\vartheta_i^c(t)}{dt} = M_d[\vartheta(t)]\, \vartheta_i^c(t) \left[F_i[\vartheta(t)] - \sum_{o=1}^{N} A_c(o) F_o[\vartheta(t)] \right] \tag{A.4}$$

Both ODE's are well defined, they have a unique solution for each initial condition and this limit does not depend on the frequency or delays in information broadcasting, as long as (11) holds. The weak convergence approach followed here allows us to interpret the learning schemes as stochastic analogs of numerical approximations of an ODE, whose r.h.s. is constructed using the Kuhn-Tucker conditions for optimality. The asymptotic behaviour of the algorithms is determined by studying the limit points $\lim_{t\to\infty} \vartheta(t)$. From (A.3) and (A.4), these are the fixed points of the equations for the conditional expected behaviour, as mentioned in Section 4.

As a final remark, this proof only requires the construction of apropriate estimators of the sensitivity that satisfy (A.2). We have provided one such method that requires basic information on moving averages, but in order to minimize information exchange, the schemes could be implemented relying on a less frequent transmission of the moving averages between nodes, yielding the same asymptotic behaviour.

Appendix B. Tabloid Method for the n-action Scheme.

One way to study the two-action scheme was introduced in [22] to study the assignment problem in terms of a tabloid solution, as follows. From the equilibrium matrix, we can consider the basic variables $x_c(i) = A_c(i) C_c \pi_i^c$ that appear in the matrix equations instead of the distribution probabilities themselves. The data of the problem in terms of the optimal throughputs can be written in the form of a tabloid, where the rows add up to the numbers in the right hand column, the columns add up to the numbers in the last row and the sum of the right hand column equals the sum of the last row. In a similar way, we can also write down the solution in terms of the variables $x_c(i)$, where the sums of the rows set equal to the corresponding number in the last column represents equation (8) and the sum of the columns set equal to the value in the last row yields equation (7).

Data						
0	λ_{10}	λ_{20}	λ_{30}	$\cdots$	λ_{N0}	C_0
λ_{01}	0	λ_{21}	λ_{31}	$\cdots$	λ_{N1}	C_1
λ_{02}	λ_{12}	0	λ_{32}	$\cdots$	λ_{N2}	C_2
λ_{03}	λ_{13}	λ_{23}	0	$\cdots$	λ_{N3}	C_3
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
λ_{0N}	λ_{1N}	λ_{2N}	λ_{3N}	$\cdots$	0	C_N
λ_0	λ_1	λ_2	λ_3	$\cdots$	λ_N	

Solution						
0	$x_0(1)$	$x_0(2)$	$x_0(3)$	$\cdots$	$x_0(N)$	C_0
$x_1(0)$	0	$x_1(2)$	$x_1(3)$	$\cdots$	$x_1(N)$	C_1
$x_2(0)$	$x_2(1)$	0	$x_2(3)$	$\cdots$	$x_2(N)$	C_2
$x_3(0)$	$x_3(1)$	$x_3(2)$	0	$\cdots$	$x_3(N)$	C_3
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$x_N(0)$	$x_N(1)$	$x_N(2)$	$x_N(3)$	$\cdots$	0	C_N
λ_0	λ_1	λ_2	λ_3	$\cdots$	λ_N	

In the solution table, some of the entries $x_c(i)$ may be zero entries, depending on the actions at that controller. If we want to find a solution for a given action matrix, we "cross" out the corresponding places in the tabloid and proceed to find a solution. It is clear from this form why the $(N-1)$-action scheme always possesses a solution of the form $x_c(i) = \lambda_{ic}$, although this solution is not unique, as shown in theorem 2.

We give now the tabloid solution of two different two-action schemes in a four node example, showing the symbol $\times$ at the zero entries where $A_c(i) = 0$.

In terms of the dimensionality of the problem, the two-action case defines a linear problem in $2N$ variables with $2N$ equations, one of them being linearly dependent on the other $2N-1$. If the problem does not decouple into several independent problems (such as two tabloids put together with zero entries in the diagonal blocks), then the solution is defined as a one dimensional subspace. This condition will later be defined in terms of equivalence classes of source nodes (see section 5). It is clear that in a four node network any two-action scheme will define a tabloid that cannot be decomposed into independent smaller tabloids, since the controllers never send permits to their own source queues.

Example 1				
0	$\lambda_1 + \lambda_0 - C_3 - a$	$C_0 + C_3 - \lambda_0 - \lambda_1 + a$	$\times$	C_0
$\times$	0	$C_1 + C_2 - \lambda_3 - a$	$\lambda_3 - C_2 + a$	C_1
a	$\times$	0	$C_2 - a$	C_2
$\lambda_0 - a$	$C_3 - \lambda_0 + a$	$\times$	0	C_3
λ_0	λ_1	λ_2	λ_3	

In order to find the one-dimensional solution set for these examples we arbitrarily choose one element of the solution matrix as a and work through the tabloid filling the entries that are determined by the requirement of the sums of rows and columns, and using $\sum_i \lambda_i = \sum_c C_c$. The general way of finding a solution in higher dimensional problems with two-action schemes that do not decouple is identical to this procedure and we omit the details of the general algorithm.

In the second example, the entry $x_1(0) = \lambda_0$ is uniquely determined by the choice of the action matrix and does not depend on the value of a. In order for this solution to be feasible in terms of the distribution probabilities, all entries of the solution matrix must be non-negative.

Example 2				
0	$\times$	$\lambda_1 + \lambda_2 - C_3 - a$	$C_0 + C_3 - \lambda_1 - \lambda_2 + a$	C_0
λ_0	0	$\times$	$C_1 - \lambda_0$	C_1
$\times$	a	0	$C_2 - a$	C_2
$\times$	$\lambda_1 - a$	$C_3 + a - \lambda_1$	0	C_3
λ_0	λ_1	λ_2	λ_3	

Since all the expressions are linear in a, the non-negativity condition can be rewritten in terms of an interval of feasibility of the form $a_0 \leq a \leq a_1$, where the limits depend on the data λ_{id}. This set may be empty, in which case the chosen action matrix will not yield the optimal solution for the original problem, therefore given a two-action scheme it does not always possess a solution. Given a problem, we may ask if it is possible to find a two-action scheme with a tabloid solution of non-negative entries. Unfortunately the answer is no, as shown in the tabloid for a counterexample, where we show the data of the problem.

Counterexample				
0	50	20	10	80
60	0	1	1	62
30	1	0	1	32
10	1	1	0	12
100	52	22	12	

In this problem there is no way that we can keep only two nonzero entries of the first row and fill out the rest of the table with non-negative entries. Although this represents an extremely unbalanced traffic where most of the throughput comes in and out of node zero, it shows that in general we cannot always find a two-action scheme which will satisfy the equilibrium equations for the optimal throughputs.

The main problem in assigning the two-action automata scheme is that the tabloids for the data and the solutions cannot be written in terms of the data of the original problem and depend on the unknown optimal throuputs. Therefore we cannot use this algorithm in order to determine if the problem admits a two-action scheme to achieve its optimum performance. This framework has further been investigated in [17] to establish the range of feasibility of the n-action design, for $n < N - 1$. This approach introduces a ficticious objective function to implement the simplex method, which finds the corner points of the feasible set, thus finding the solution points of the tabloids.

Appendix C. Notation.

N : number of nodes in the network

W : number of permits in the network

C_{ij} : capacity of link (i, j)

D_{ij} : distance between node i and node j

r_{ij}^d : probability of routing a packet with destination d from i to j

p_{ij} : $= D_{ij}/c$ propagation delay at trunks

p_j^d : propagation delay from controller d to source node j

π_i^d : control variable: probability of sending a permit from d to source node i

s_i : relative number of visits to permit source queue i

c_j : relative number of visits to controller queue j

f_{jl} : relative number of visits to trunk queue (j, l)

Λ_{id} : external arrival rate at node i of packets with destination d

α : average packet size

λ_{id} : stationary average throughput of packets with origin-destination (i, d)

λ_i : aggegate throughput $= \sum_d \lambda_{id}$

C_d : aggegate throughput $= \sum_i \lambda_{id}$

T_{id} : stationary average end-to-end delay of packets with origin-destination (i, d)

$A_c(i)$: action matrix $= \mathbf{1}_{\{\pi_i^c > 0\}}$