TECHNICAL RESEARCH REPORT

Convergence Properties For Uniform Ant Routing

by Joon-Hyuk Yoo, Richard J. La and Armand M. Makowski

CSHCN TR 2003-15
(ISR TR 2003-24)

The Center for Satellite and Hybrid Communication Networks is a NASA-sponsored Commercial Space Center also supported by the Department of Defense (DOD), industry, the State of Maryland, the University of Maryland and the Institute for Systems Research. This document is a technical report in the CSHCN series originating at the University of Maryland.

Web site http://www.isr.umd.edu/CSHCN/

Convergence Properties For Uniform Ant Routing

Joon-Hyuk Yoo, Richard J. La and Armand M. Makowski

The authors are with the Department of Electrical and Computer Engineering, and the Institute for Systems Research, University of Maryland, College Park, MD 20742. E-mail: juneki@glue.umd.edu, hyongla@eng.umd.edu and armand@isr.umd.edu.

Abstract

We study the convergence property of a family of distributed routing algorithms based on the ant colony metaphor, which generalize the uniform ant routing algorithms proposed earlier. For a simple two-node network, we show that the probabilistic routing tables under these algorithms converge in distribution, and discuss some of implementation issues.

I. INTRODUCTION

In the past decade, several authors have used the ant colony metaphor to design distributed adaptive routing algorithms in both datagram networks [4] and telephone networks [5]; a good survey of these efforts is given in Chapter 2 of [1]. In this paper we revisit the class of *uniform ant routing* algorithms discussed by Subramanian et al. [6]. These algorithms are randomized algorithms that implement a form of *backward* learning in response to short control messages called *ants*. This can be loosely described as follows:

Consider the situation where hosts are provided connectivity through a network of R routers. Two routers are said to be neighbors if there exists a bidirectional point-to-point link between them, and let $\mathcal{N}_r$ denote the set of routers which are neighbors of router $r = 1, \ldots, R$. Router r maintains a probabilistic routing table with a separate vector entry $(d, (i, p_i),\ i \in \mathcal{N}_r)$ for each host destination d. For each neighboring router i in $\mathcal{N}_r$, we understand p_i as the probability with which router r uses link (r, i) when forwarding a *data* packet destined for d. There is a cost c_{ri} associated with the use of link (r, i); this cost is assumed symmetric (i.e., $c_{ri} = c_{ir}$) and is known to router r.

An ant is a control message of the form $(d, s\|c)$ where d and s are distinct hosts, and c is some numerical value to be updated in due course. We refer to hosts d and s as the destination and source, respectively, and regard c as an *estimate* of the cost-to-go for reaching host d. Periodically, host d generates an ant $(d, s\|c)$ which is destined for some randomly selected host $s \neq d$ with c initially set to zero. The ant is forwarded to the source s over the network of interconnected routers and on the way updates the routing tables at intermediary routers (in a way to specified shortly).

When ant $(d, s\|c)$ arrives at the intermediary router r coming from router i through link (i, r), the cost estimate c for reaching d from router i is incremented by the cost of the (reverse) link (r, i) with

$$c \leftarrow c + c_{ri} \tag{1}$$

and the new value of c thus provides an estimate of the cost-to-go to reach d from router r. Next, the vector entry in the routing table for destination d is updated according to

$$p_i \leftarrow \frac{p_i + \Delta p}{1 + \Delta p}, \quad p_j \leftarrow \frac{p_j}{1 + \Delta p}\ \ (j \in \mathcal{N}_r \setminus \{i\}) \tag{2}$$

where $\Delta p = \frac{K}{f(c)}$ $(K > 0)$, and $f(c)$ is a non-decreasing function of the just incremented value of c. Consequently, the probability of using the (reverse) link over which the ant arrived at router r has been increased relative to that of other links, while that of other links is being discounted. This constitutes a form of (backward) *reinforcement learning*.

Upon completing the updates (1)-(2), router r forwards the updated ant $(d, s\|c)$ to one of its neighboring router i in $\mathcal{N}_r$. The ant $(d, s\|c)$ eventually reaches its destination s with c now giving the end-to-end cost of sending a message from s to d (in fact the cost of the very path followed by the ant), and is destroyed. The manner in which ant forwarding is carried out distinguishes the various types of ant routing algorithms. The *uniform* ant algorithm is designed with multi-path routing in mind, and requires that ants arriving, say at router r, be forwarded to the next neighboring router with equal probability L_r^{-1} where L_r is the number of routers in $\mathcal{N}_r$, and more generally, the number of links going out of r (in the case of multiple links). As a result, uniform ants will utilize each and every path between a source and a destination with a positive probability.

In [6], Subramanian et al. discuss among other things the convergence of uniform ant algorithms. Their discussion focuses on a simple two router network for which the algorithm is claimed [6, Prop. 2] [7] to converge (in an unspecified mode of convergence) to constant values. Here we revisit this two-node model. Our main result is contained in Theorem 1, and states that (i) convergence takes place in *distribution*, and (ii) *not* to constants. This last fact has implications for the implementation of such ant algorithms.

The key observation behind the proof of Theorem 1 is the identification of the iterates of the uniform ant algorithm with the output of a very simple collection of iterated random functions. A large literature is available on the convergence of these iterative schemes, and the survey in [3] (and references therein) provides a nice introduction to this topic. We exploit the very simple structure of the underlying collection of random functions to give a simple and self-contained proof of Theorem 1.

The paper is organized as follows: The uniform ant algorithm is formally described in Section II for a simple two-link network. Theorem 1 is presented in Section III. A convergence proof is given in Section IV, and implementation issues are pointed out in Section V.

II. A TWO-NODE NETWORK WITH PARALLEL LINKS

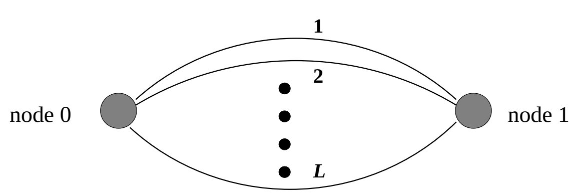

Fig. 1. Parallel link network.

In this section we present a two-node model together with the corresponding uniform ant routing algorithms. The setup is somewhat more general than the one used in [6].

The network comprises two routers or nodes, thereafter labeled node $i = 0$ and node $i = 1$, connected by a set of L parallel bidirectional links (Fig. 1). All hosts are attached to either node $i = 0$ or node $i = 1$.

Pick node i ($i = 0, 1$). Destinations hosts attached to node i generate ants at times $\{t_n^i,\ n = 1, 2, \ldots\}$ with $0 < t_n^i < t_{n+1}^i$ for each $n = 1, 2, \ldots$. Forwarding the n^{th} ant to node $1 - i$ requires that one of the L links from node i to node $1-i$, say $\ell_{1-i}(n)$, be selected. The n^{th} ant is then sent over link $\ell_{1-i}(n)$, and arrives at node $1-i$, say at time a_n^i with $t_n^i < a_n^i$. For simplicity of exposition we assume the non-overtaking condition

$$a_n^i < a_{n+1}^i, \quad n = 1, 2, \ldots \tag{3}$$

with $a_0^i = 0$ for the sake of concreteness.

Node i maintains a probabilistic forwarding table $\mathbf{p}^i(n) := (p_1^i(n), \cdots, p_L^i(n))$ with $0 \leq p_\ell^i(n) \leq 1$ ($\ell = 1, \ldots, L$) and $\sum_{\ell=1}^L p_\ell^i(n) = 1$. The entry $p_\ell^i(n)$ is interpreted as the probability that during the interval $[a_n^{1-i}, a_{n+1}^{1-i})$, node i forwards a data packet from node i to node $1-i$ over link ℓ.

When at time a_n^i, node $1-i$ receives the n^{th} ant, say over link $\ell_{i-1}(n) = \ell$, it immediately updates its probabilistic forwarding table according to

$$p_\ell^{1-i}(n+1) = \frac{p_\ell^{1-i}(n) + C_\ell(1-i)}{1 + C_\ell(1-i)} \tag{4}$$

and

$$p_k^{1-i}(n+1) = \frac{p_k^{1-i}(n)}{1 + C_\ell(1-i)}, \quad k \neq \ell, \ k = 1, \ldots, L \tag{5}$$

for constants $C_\ell(1-i) > 0$. The selection of these constants is discussed later. The cost update (1) is superfluous here due to the simplified structure of this two-node network.

To specify how ants are propagated, we shall assume throughout the following : The $\{1, \ldots, L\}$-valued random variables (rvs) $\{\ell_0(n), \ell_1(n), \ n = 1, 2, \ldots\}$ are *mutually independent* rvs which are taken to be *independent* of the initial conditions $\mathbf{p}^0(0)$ and $\mathbf{p}^1(0)$. Moreover, for each $i = 0, 1$, the rvs $\{\ell_i(n), \ n = 1, 2, \ldots\}$ are *i.i.d.* rvs distributed according to some probability mass function (pmf) $\mathbf{v}^i = (v_1^i, \ldots, v_L^i)$ on $\{1, \ldots, L\}$.

III. The convergence result

Under the assumptions made above, it is plain from (4) and (5) that the probabilistic forwarding tables of the nodes are updated independently of each other, whence the evolutions of the tables $\{\mathbf{p}^0(n), \ n = 0, 1, \ldots\}$ and $\{\mathbf{p}^1(n), \ n = 0, 1, \ldots\}$ are decoupled. To describe the long-term behavior of these forwarding tables, we find it convenient to introduce the following notation: For any integer L, the ℓ^{th} component of any element $\mathbf{x}$ in $\mathbb{R}^L$ is denoted either by x^ℓ or by x_ℓ, $\ell = 1, \ldots, L$, so that $\mathbf{x} \equiv (x^1, \ldots, x^L)$ or $(x_1, \ldots, x_L)$. A similar convention is used for $\mathbb{R}^L$–valued rvs.

Fix $i = 0, 1$, and for each link $\ell = 1, \cdots, L$, define the mapping $\phi_\ell^i : [0,1]^L \to [0,1]^L$ by

$$\phi_{\ell,k}^i(\mathbf{p}) := \begin{cases} \frac{p_\ell + C_\ell(i)}{1 + C_\ell(i)} & k = \ell \\ \frac{p_k}{1 + C_\ell(i)} & k \neq \ell \end{cases}, \quad \mathbf{p} \in [0,1]^L \tag{6}$$

for constants $C_\ell(i) > 0$.

Theorem 1: Under the foregoing assumptions we have:

(i) For each $i = 0, 1$, *the limit*

$$\mathbf{p}_\star^i = \lim_{n \to \infty} \left(\phi_{\ell_i(1)}^i \circ \phi_{\ell_i(2)}^i \circ \ldots \circ \phi_{\ell_i(n)}^i \right)(\mathbf{p}) \tag{7}$$

exists for each $\mathbf{p}$ *in* $[0,1]^L$*, and is independent of* $\mathbf{p}$*;*

(ii) Moreover, there exists a pair of independent $[0,1]^L$*-valued rvs* $\mathbf{p}^0$ *and* $\mathbf{p}^1$ *distributed like* $\mathbf{p}^0_\star$ *and* $\mathbf{p}^1_\star$*, respectively, such that*

$$(\mathbf{p}^0(n), \mathbf{p}^1(n)) \Longrightarrow_n (\mathbf{p}^0, \mathbf{p}^1) \tag{8}$$

with $\Longrightarrow_n$ *denoting convergence in distribution (or in law) (with* n *going to infinity).*

The proof of Theorem 1 is given in Section IV. We emphasize that the result holds for a class of algorithms which is somewhat more general than the one introduced in [6]: Indeed, the constants entering (4)-(5) are arbitrary and need not be constrained to

$$C_\ell(1-i) = \frac{K}{f(c_\ell)} =: \Delta_\ell \quad i = 0,1,\ \ell = 1,\ldots,L \tag{9}$$

with cost c_ℓ for using link ℓ in either direction, constant $K > 0$ and a strictly increasing function $f : \mathbb{R} \to (0,\infty)$. Next, when the pmfs $\mathbf{v}^0$ and $\mathbf{v}^1$ are assumed to be the uniform pmf on $\{1,\ldots,L\}$, we recover the case discussed in [6], hence the name uniform ant algorithm.

Finally, we note that the assumptions enforced on the "reception" times $\{a^i_n,\ n = 1,2,\ldots\}$ $(i = 1,0)$ can be weakened considerably: These times could in principle be *random* and Theorem 1 would still hold provided they are assumed independent of the link selections rvs $\{\ell_0(n), \ell_1(n),\ n = 1,2,\ldots\}$. The non-overtaking assumption (3) can also be dropped if we now interpret $\ell_{1-i}(n)$ as the identity of the link from node i to node $1-i$ which was traversed by the n^{th} ant *received* at node $1-i$ at the (possibly random) time a^i_n. Then, under the aforementionned independence assumptions, Theorem 1 will still hold.

Subramanian et al. [6, Prop. 2] claim the convergence

$$\lim_{n\to\infty} \mathbf{p}^i(n) = L_i, \quad i = 0,1$$

for some *constants* L_0 and L_1, without further indication of the mode of convergence used for this convergence statement which involve rvs. The validity of these claims is further dispelled in Section V.

IV. A PROOF OF THEOREM 1

As remarked earlier, because the forwarding probability tables evolve independently of each other, we need only establish the one-dimensional convergence $\mathbf{p}^i(n) \Longrightarrow_n \mathbf{p}^i_\star$ for each $i = 0,1$.

Thus fix $i = 0,1$. With the notation above, the updating rules (4) and (5) can be written more compactly as

$$\mathbf{p}^i(n+1) = \phi^i_\ell(\mathbf{p}^i(n)) \quad \text{if} \quad \ell_i(n+1) = \ell, \quad n = 0,1,\ldots$$

with $\mathbf{p}^i(0)$ denoting the forwarding probability vector initially stored at node i (i.e., at time $a^i_0 = 0$).

Fix $n = 0,1,\ldots$. Iterating yields the relation

$$\begin{aligned} &\mathbf{p}^i(n+1) \\ = \quad &\left(\phi^i_{\ell_i(n+1)} \circ \phi^i_{\ell_i(n)} \circ \ldots \circ \circ\phi^i_{\ell_i(2)} \circ \phi^i_{\ell_i(1)}\right)(\mathbf{p}^i(0)). \end{aligned} \tag{10}$$

The key observation is the stochastic equivalence [1]

$$\begin{aligned} & \mathbf{p}^i(n+1) \\ =_{st} \quad & \left(\phi^i_{\ell_i(1)} \circ \phi^i_{\ell_i(2)} \circ \ldots \circ \phi^i_{\ell_i(n)} \circ \phi^i_{\ell_i(n+1)}\right)(\mathbf{p}^0(0)) \end{aligned} \tag{11}$$

which holds by virtue of the i.i.d. assumption on the sequence $\{\ell_i(n),\ n = 1, 2, \ldots\}$ and its independence from $\mathbf{p}^0(0)$. The one-dimensional convergence $\mathbf{p}^i(n) \Longrightarrow_n \mathbf{p}_\star(n)$ will follow if we can show the *pointwise* convergence

$$\lim_{n\to\infty} \mathbf{p}^i(n) = \mathbf{p}^i_\star. \tag{12}$$

To do so, we equip $\mathbb{R}^L$ with the norm defined by $\|\mathbf{x}\| := \sum_{\ell=1}^L |x_\ell|$ for any vector $\mathbf{x}$ in $\mathbb{R}^L$. This norm is equivalent to the usual Euclidean norm, but easier to use here.

Fix $\ell = 1, \ldots, L$. For arbitrary $\mathbf{x}$ and $\mathbf{y}$ in $[0,1]^L$, we get

$$\begin{aligned} & \|\phi^i_\ell(\mathbf{x}) - \phi^i_\ell(\mathbf{y})\| \\ = \quad & \sum_{k=1}^L \left|\phi^i_{\ell,k}(\mathbf{x}) - \phi^i_{\ell,k}(\mathbf{y})\right| \\ = \quad & \sum_{k\neq\ell} \left|\frac{x_k - y_k}{1 + C_\ell(i)}\right| + \frac{|x_\ell + C_\ell(i) - (y_\ell + C_\ell(i))|}{1 + C_\ell(i)} \\ = \quad & \sum_{k=1}^L \frac{|x_k - y_k|}{1 + C_\ell(i)} \\ = \quad & \frac{\|\mathbf{x} - \mathbf{y}\|}{1 + C_\ell(i)} \\ \leq \quad & K_i \|\mathbf{x} - \mathbf{y}\| \end{aligned}$$

with

$$K_i := \max_{\ell=1,\ldots,L} (1 + C_\ell(i))^{-1} < 1.$$

Therefore,

$$\max_{\ell=1,\ldots,L} \|\phi^i_\ell(\mathbf{x}) - \phi^i_\ell(\mathbf{y})\| \leq K_i \|\mathbf{x} - \mathbf{y}\| \ . \tag{13}$$

Using this last inequality, it is a simple matter to check (by induction on $n = 1, 2, \ldots$) that

$$\begin{aligned} & \| \left(\phi^i_{\ell_i(1)} \circ \ldots \circ \phi^i_{\ell_i(n-1)} \circ \phi^i_{\ell_i(n)}\right)(\mathbf{p}) \\ & \qquad - \left(\phi^i_{\ell_i(1)} \circ \ldots \circ \phi^i_{\ell_i(n-1)} \circ \phi^i_{\ell_i(n)}\right)(\mathbf{q}) \| \\ \leq \quad & K_i^n \|\mathbf{p} - \mathbf{q}\| \end{aligned} \tag{14}$$

for arbitrary $\mathbf{p}$ and $\mathbf{q}$ in $[0,1]^L$. Furthermore, for each $m = 1, 2, \ldots$, we get

$$\| \left(\phi^i_{\ell_i(1)} \circ \phi^i_{\ell_i(2)} \circ \ldots \circ \phi^i_{\ell_i(n+m)}\right)(\mathbf{p})$$

[1] Two $\mathbb{R}$–valued rvs X and Y are said to be equal in law if they have the same distribution, a fact we denote by $X =_{st} Y$.

$$
\begin{aligned}
& - \left(\phi^i_{\ell_i(1)} \circ \phi^i_{\ell_i(2)} \circ \ldots \circ \phi^i_{\ell_i(n)}\right)(\mathbf{p})\| \\
\leq \quad & K_i^n \| \left(\phi^i_{\ell_i(n+1)} \circ \phi^i_{\ell_i(n+2)} \circ \ldots \circ \phi^i_{\ell_i(n+m)}\right)(\mathbf{p}) - \mathbf{p}\| \\
\leq \quad & K_i^n L \quad \text{uniformly in } m. \qquad (15)
\end{aligned}
$$

From the fact $K_i < 1$, it follows for each $\mathbf{p}$ that the sequence $\{ \left(\phi^i_{\ell_i(1)} \circ \phi^i_{\ell_i(1)} \circ \ldots \circ \phi^i_{\ell_i(n)}\right)(\mathbf{p}), \quad n = 1, 2, \ldots\}$ forms a Cauchy sequence, hence is convergent. Eqn. (14) shows that the limit of this convergent sequence is independent of $\mathbf{p}$. This establishes (12) and the proof is now complete.

V. Implementation issues

In this section we briefly discuss some of the implementation issues associated with uniform ant routing algorithms. Consider the case first described in [6] with $L = 2$ where the constants in the probability updates (4) and (5) are selected according to (9) and for $i = 0, 1$, the pmf $\mathbf{v}^i$ is the uniform pmf on $\{1, 2\}$. It is easy to see that the probabilistic routing tables evolve according to

$$
\mathbf{p}^i(n+1) = \begin{cases} \left[\frac{p_1^i(n)+\Delta_1}{1+\Delta_1}, \frac{p_2^i(n)}{1+\Delta_1}\right] & \text{w.p. } \frac{1}{2} \\ \\ \left[\frac{p_1^i(n)}{1+\Delta_2}, \frac{p_2^i(n)+\Delta_2}{1+\Delta_2}\right] & \text{w.p. } \frac{1}{2} \end{cases} \qquad (16)
$$

where the constants Δ_ℓ ($\ell = 1, 2$) are given by (9).

Selecting the proper values of K and $f(c_\ell)$ ($\ell = 1, 2$) is crucial here for good performance. Indeed, if the parameters are selected such that

$$
\frac{1}{1+\Delta_2} < \frac{\Delta_1}{1+\Delta_1},
$$

then the iterates produced by (16) exhibit an oscillatory behavior with successive values possibly bouncing around between the *non-overlapping* intervals $(0, (1+\Delta_2)^{-1})$ and $((1+\Delta_1)^{-1}\Delta_1, 1)$. This oscillatory behavior leads to undesirable oscillations in the routing tables, and is further evidence that the convergence of Theorem 1 cannot be in the a.s. sense. In fact, convergence takes place in distribution (*not* a.s.) to a limiting random variable whose distribution has a *non-connected* support on the interval $[0, 1]$.

References

[1] E. Bonabeau, M. Dorigo and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Santa Fe Institute Studies in the Sciences of Complexity, Oxford University Press, New York, N.Y., 1999.

[2] P. Billingsley, *Convergence of Probability Measures*, John Wiley & Sons, New York (NY), 1968.

[3] P. Diaconis and D. Freedman, “Iterated random functions,” *SIAM Review* **41** (1999), pp. 45-76.

[4] G. DiCaro and M. Dorigo, “Two ant colony algorithms for best-effort routing in datagram networks,” in Proceedings of PDCS 1998: International Conference on Parallel and Distributed Computing and Systems, Anaheim (CA).

[5] R. Schoonderwoerd, O. Holland, J. Bruten and L. Rothkrantz, “Ant-based load balancing in telecommunications networks,” *Adaptive Behavior* **5** (1996), pp. 169-207.

[6] D. Subramanian, P. Druschel and J. Chen, “Ants and reinforcement learning: A case study in routing in dynamic networks,” in Proceedings of IJCAI 1997: The International Joint Conference on Artificial Intelligence, Nagoya (Japan), August 1997.

[7] D. Subramanian, P. Druschel and J. Chen, “Ants and reinforcement learning: A case study in routing in dynamic networks,” Technical Report **TR96-259**, Department of Computer Science, Rice University (updated July 1998).