Computer Systems

G. Bell, D. Siewiorek, and S. H. Fuller, Editors

On Quadratic Adaptive Routing Algorithms

Carson E. Agnew
Stanford University

Two analytic models of a store-and-forward communications network are constructed, one to find the optimal message routing and the other to illustrate the equilibrium (stationary state) maintained by an adaptive routing algorithm. These models show that adaptive routing does not satisfy the necessary conditions for an optimal routing. Adaptive routing tends to overuse the direct path and underuse alternate routes because it does not consider the impact of its current routing decision on the future state of the network. The form of the optimality conditions suggests that a modification of the adaptive algorithm will result in optimality. The modification requires the substitution of a quadratic bias term instead of a linear one in the routing table maintained at each network node. Simulation results are presented which confirm the theoretical analysis for a simple network.

Key Words and Phrases: routing algorithms, adaptive routing, quadratic routing, alternate routing, store-and-forward network, distributed network, computer network, message switching

CR Categories: 3.81, 8.0, 8.1, 8.3

1. Introduction

Communications networks which use a store-and-forward strategy for transmitting messages must choose the path along which to relay the message. Routing algorithms for such networks are decision rules which attempt to minimize the average time a message spends in the network given a fixed network topology and a destination address. Central and local control methods and fixed and adapting routing schemes have all been tried in actual networks or studied by simulation methods [5, 11]. In addition, several studies have proposed solution algorithms for the "optimal routing" problem where minimum transmission time routes are found simultaneously for all message traffic [3, 4, 11].

One can conclude from these studies that alternate routing, which sends some traffic over the shortest source-to-destination path and routes the rest by a more roundabout path, is generally superior to the use of a single route. Also, adaptive algorithms, which use information about the actual delays in the network to estimate the current shortest path, have been found to be superior to fixed routing schemes in which the same route (or routes) is used regardless of the actual state of the network.

Because of this superiority and because they are relatively easy to implement, adaptive routing schemes using alternate routing are of considerable interest. However, this paper will show that such algorithms do not automatically yield an optimal routing. This will be seen by comparing the necessary conditions for an optimal routing with the equilibrium conditions which an adaptive routing scheme satisfies in a simple analytic model. The comparison will also suggest a simple modification to the routing algorithm which makes the adaptive scheme satisfy the optimality conditions and leads to reduced delays for simple networks.

2. Optimal Alternate Routing

Consider any particular source node and destination node, with the n routes between them indexed by $i = 1, 2, \cdots, n$. The simplest optimal routing problem involves distributing the traffic at the input node over the available routes so as to minimize the average delay from the time the message enters the source until it arrives at the destination.

In this model it is assumed that messages arrive in a Poisson manner at the source node with rate γ per second and have exponentially distributed lengths of $1/\mu$ bits. Messages are queued in one of n buffers corresponding to the outgoing routes. Each buffer is assumed to have an infinite capacity. On each route the first link has a channel capacity of C_i bps, $i = 1, 2, \ldots, n$, so that the time spent in the source node by a message assigned to the ith route is T_i seconds, with expected value:

$$\bar{T}_i = 1/(\mu C_i - \lambda_i) \tag{1}$$

where λ_i is the rate at which traffic is assigned to the ith route in messages per second.

In addition to T_i, messages taking the ith route incur an additional delay T_i' after they leave the source node. T_i' includes propagation delays and processing and queuing delays at nodes encountered along route i

This research was supported by the National Science Foundation Grant GJ 36392X. Author's present address: Mathematica, Inc., P.O. Box 2392, Princeton, NJ 08540.

between the source and the sink, so the sum $T_i + T_i'$ is the network transmission time over the ith route. The average transmission delay to all messages in the network is given by the well-known formula [6]:

$$\bar{T} = (1/\gamma) \sum_{i=1}^{n} \lambda_i (T_i + T_i'). \tag{2}$$

The optimal routing problem is that of minimizing $\bar{T}$, subject to the constraints

$$\gamma = \sum_{i=1}^{n} \lambda_i, \; 0 \le \lambda_i \le \mu C_i, \; i = 1, 2, \ldots, n. \tag{3}$$

The problem can be solved in practice by numerical methods (see [4, 11], but since this paper is concerned with the theoretical implications of adaptive routing, the Kuhn-Tucker theorem [7] will be used to study the conditions which an optimal routing must satisfy. Forming the Lagrangian from (1)–(3) gives:

$$\begin{aligned} \mathcal{L} = {} & (1/\gamma) \sum_{i=1}^{n} (\lambda_i/(\mu C_i - \lambda_i) + \lambda_i T_i') \\ & - p \sum_{i=1}^{n} \lambda_i + \sum_{i=1}^{n} r_i(\lambda_i - \mu C_i) - \sum_{i=1}^{n} q_i \lambda_i \end{aligned} \tag{4}$$

where p, r_i, and q_i are undetermined multipliers. The conditions satisfied by our optimal choice of the λ_i are:

$$((\mu C_i/\gamma)/(\mu C_i - \lambda_i)^2) + T_i'/\gamma - p + r_i - q_i = 0, \quad i = 1, 2, \ldots, n \tag{5}$$

$$r_i(\lambda_i - \mu C_i) = 0, \; r_i \ge 0, \quad i = 1, 2, \ldots, n \tag{6}$$

$$q_i \lambda_i = 0, \qquad q_i \ge 0, \quad i = 1, 2, \ldots, n \tag{7}$$

as well as the constraints (3). Unless the source is completely saturated (i.e. $\gamma \ge \sum_i \mu C_i$), it is easy to see that $\lambda_i < \mu C_i$ holds for all i because an optimal routing must keep any link from saturating before any other link. This implies $r_i = 0$ for all $i = 1, 2, \ldots, n$.

The condition on q_i implies that $q_i = 0$ if any traffic is assigned to the ith route ($\lambda_i > 0$) and, conversely, that if no traffic has been assigned to the ith route then $q_i \ge 0$. Assuming non-saturation of the source we may restate the necessary conditions without p or q_i as

$$\begin{aligned} & (\mu C_i/(\mu C_i - \lambda_i)^2) + T_i' \\ & \quad = (\mu C_j/(\mu C_j - \lambda_j)^2) + T_j' \text{ if} \\ & \qquad \lambda_i > 0, \; \lambda_j > 0 \\ & (\mu C_i/(\mu C_i - \lambda_i)^2) + T_i' \\ & \quad \le (\mu C_j/(\mu C_j - \lambda_j)^2) + T_j' \text{ if} \\ & \qquad \lambda_i > 0, \; \lambda_j = 0 \end{aligned} \tag{8}$$

3. Equilibrium Traffic Assignments

Now let us consider a model of an adaptive routing algorithm in order to contrast its performance with that of an optimal routing scheme. This model is typical of routing schemes used in actual computer networks (as discussed in [5]) in which no global routing procedure is used. Instead, each node in the network maintains a table with an entry for each possible destination and each neighboring node. The entry in the table for location (i, j) is the estimated time required for a message sent via the ith neighbor to reach the jth destination, including the time spent in the sending node awaiting transmission. The table is kept current by having neighbors periodically exchange their estimates of the minimum time to reach each destination and using this information to update each node's routing table. When there is only one destination, as in Section 2, the routing table at the sending node is a vector whose ith element is T_i' plus a linear function giving an estimate of the waiting time in the buffer $T_i = (1 + L_i)/\mu C_i$ where there are L_i messages in the buffer (including the one in service).

The table is used to route messages in the following manner. Upon arrival at the sending node receipt of the message is acknowledged and the node consults the message's destination address. If it is addressed to the jth destination, the node then consults column j of its routing table, chooses the minimum entry in the column and places the message in the queue for the associated transmission line. Finally, the node updates the routing table to reflect the new state of the queue. (Updating will also occur after each transmitted message is acknowledged.) In our simple model, we may think of the node as assigning the message to the route i for which the current value of $T_i + T_i'$ is a minimum.

Invoking now all of the assumptions made in finding the optimal routing, namely Poisson arrivals and exponential message lengths, we find that it is almost intuitively obvious that the transmission times along the various routes must satisfy the average conditions

$$\begin{aligned} & 1/(\mu C_i - \lambda_i) + T_i' = 1/(\mu C_j - \lambda_j) + T_j' \\ & \qquad \text{if } \lambda_i > 0 \text{ and } \lambda_j > 0 \\ & 1/(\mu C_i - \lambda_i) + T_i' < 1/(\mu C_j - \lambda_j) + T_j' \\ & \qquad \text{if } \lambda_i > 0 \text{ and } \lambda_j = 0. \end{aligned} \tag{9}$$

That is, the expected time for a message to travel from the source to the destination node is the same over any route to which traffic is assigned and is less than the travel time over any route which has no traffic assigned to it.

To see why these conditions must hold, suppose there were two routes (say 1 and 2) with nonzero traffic λ_1 and λ_2 but with inequality holding among the expected message delays. But this inequality cannot be true of the expected values of the message delays. For we know that if it were true that adaptive routing algorithm would assign all arriving messages to the one route with least delay, either bringing the expected delay on the two routes into equality or making λ_1 or λ_2 zero.

The most interesting implication of this equilibrium property of adaptive routing is that it does not correspond to the necessary conditions satisfied by an optimal routing (8). Thus the traffic routing provided by an adaptive routing algorithm with a linear bias cannot be optimal.[1] This is a rather surprising result, since it would seem that an algorithm which routes each message so as to minimize its transmission delay ought also to minimize the overall average delay per message. The reason why adaptive routing does not attain the overall optimum is that each individual minimization based on the routing table neglects the effect which the current decision has on future choices. Placing a message in the queue for any particular output channel imposes an additional delay on subsequent messages which arrive before the queue dissipates. As a result, subsequent messages must either join the queue and wait an additional length of time in this node for transmission or select an alternate route which entails extra delay somewhere else in the network. This implicit cost is not taken into account by an adaptive routing algorithm because the estimated single-message delay which is entered in the routing table does not include any penalty which reflects the effect of the current routing choice on the delays incurred by subsequent messages.

4. An Instructive Example

Before suggesting a way to make adaptive routing algorithms achieve the minimum average delay, a simple two-route example will be considered. In it the arriving traffic may be assigned either to the "direct" route from source to destination (numbered 1) or the "indirect" or "alternate" route (numbered 2). Both routes have equal capacity: $\mu C_1 = \mu C_2 = 1$. On the direct route $T_1' = 0$ (there is no additional delay) while on the indirect a message requires $T_2' = 1.0$ seconds extra. As before, the total input traffic is assumed to be Poisson, arriving at rate γ, with exponential message lengths, and the buffers for routes 1 and 2 have infinite capacity.

Table I shows both the traffic assignment which is found by adaptive routing and the optimal assignment for various values of γ. From this table we can draw two important conclusions which help us understand the difference between optimal and adaptive routing.

First, we can see that the optimal assignment routes more traffic over the alternate route and initiates alternate routing at lower γ than does the adaptive algorithm. In general adaptive routing tends to overuse the direct route. This is precisely what the myopic nature of adaptive routing should lead one to expect. For if each assignment of a message neglects the possibility that subsequent messages will be delayed, more traffic than is optimal will be assigned to the direct route, which appears to give the shorter delay.

Second, Table I indicates that the difference between the optimal T and the T given by adaptive routing is small when the traffic is either very light or very heavy, and is greatest when the system is moderately loaded. The reason for this is also intuitive. Under light traffic conditions alternate routing in any form is unattractive because the average queuing delay on the direct route is much lower than on the indirect route. Therefore neither optimal nor adaptive routing will assign any traffic to the indirect route. Similarly, when the network is almost saturated the queuing delays are so long over each route that neither optimal nor adaptive routing can do anything but avoid saturating any one link before any other. Hence traffic must be assigned almost equally to each route. The superiority of optimal routing at moderate loads is important because most networks are designed to operate at just these traffic levels. The use of an adaptive scheme like the one discussed here can thus result in inefficient operation of an otherwise well-designed network.

Table I. Comparison of Optimal and Adaptive Routing.

	Assignment adaptive routing			Optimal assignment		
Source traffic rate γ	Direct route	Alternate route	$\bar{T}$	Direct route	Alternate route	$\bar{T}$
0.20	0.200	0.000	1.250	0.200	0.000	1.250
0.50	0.500	0.000	2.000	0.352	0.148	1.730
1.00	0.618	0.382	2.618	0.531	0.469	2.484
1.50	0.781	0.719	4.562	0.754	0.746	4.499
1.80	0.905	0.895	10.52	0.900	0.900	10.50

5. Improving Adaptive Routing Algorithms

Because adaptive algorithms are flexible, easy to implement, and have other features which are useful in a distributed network, it would be convenient if a way could be found to make them route messages so as to actually minimize the average transmission delay. Fortunately, an easy modification can be made which makes the adaptive algorithm also an optimal one.

The rationale for the modification can be seen by comparing the necessary conditions for an optimal route (8) to the equilibrium conditions (9). Notice that if a quantity whose expected value is $\mu C_i/(\mu C_i - \lambda_i)^2$ were added to T_i' each time the routing table is updated these two conditions would be identical and the steady state achieved by an adaptive routing algorithm would satisfy the necessary conditions for optimality. In fact, it is possible to find this quantity:

$$S_i \triangleq (1/\mu C_i)\,(1 + L_i)(1 + L_i/2) \tag{10}$$

which is a quadratic function of the queue length L_i. The equilibrium conditions are now

$$\begin{aligned} \bar{S}_i + T_i' &= \bar{S}_j + T_j' \quad \text{if } \lambda_i > 0 \text{ and } \lambda_j > 0 \\ \bar{S}_i + T_i' &< \bar{S}_j + T_j' \quad \text{if } \lambda_i > 0 \text{ and } \lambda_j = 0 \end{aligned} \tag{11}$$

[1] It is interesting to note that the conditions (9) are optimal for an extreme case of an objective like that proposed by Meister, Mueller and Rudin [9]. This objective seeks to minimize $T^{(k)} = [\sum_{i=1}^{n} (\lambda_i/\gamma)(\bar{T}_i + T_i')^k]^{1/k}$. When $k \to \infty$ the optimal assignment makes the expected transmission time the same on all routes.

and correspond to the necessary conditions for our optimal assignment. Thus the equilibrium routing is also optimal.

This scheme for updating the routing table, which we call "quadratic routing" because of the functional form of (10), corresponds to the use of an estimate of the incremental delay (a measure of the gradient of the objective function) instead of the average delay in making individual routing decisions [4]. The procedure works because using S_i instead of the linear function T_i will correct the differences between the optimal and adaptive routing assignments by penalizing assignments to routes with long queues. The quadratic term will be larger on the average where the buffer is fuller, i.e. on the direct route. As a result the direct route will be used less and the alternate route more.

From a practical standpoint the quadratic routing algorithm requires only one multiplication, one addition, and one division by two more per table entry than the old adaptive routing algorithm; hence the increase in computational overhead is slight. No more information is required before each node can compute the value of S_i, so there need be no increase at all in overhead messages exchanged between nodes. When the source node sends out information for routing at other nodes, it now transmits the current $T_i' + S_i$ instead of $T_i' + T_i$, so no programming changes should be required in this portion of the adaptive algorithm. Finally, since the same procedure is used for routing individual messages as before, all this code could be kept intact too.

There are several objections which might nevertheless make it impractical to implement quadratic routing in an actual network. First, it is based on a very simple and abstract model. In Section 6, however, evidence is presented which shows the algorithm to be more robust than the simple theory might lead one to believe. Second, since the model is not dynamic, the optimality conditions which are enforced by quadratic routing may not be optimal except when the equilibrium is attained. Reference [1] shows, however, that a good suboptimal dynamic control guides the system toward the optimal equilibrium point (8) and keeps it in that neighborhood. Hence quadratic routing should be nearly optimal.

A third objection is that the models presented here assume that the additional delay to a message after it leaves the source node, T_i', is known when in real networks the nodes know only estimates of T_i' which are updated relatively infrequently. Thus the apparently optimal solution given by the model may not be optimal in a real system. Whether this is so is an important empirical question which cannot be answered here. However, it seems reasonable that the relative superiority of quadratic routing should be insensitive to changes in the quality of the information possessed by the nodes because both routing schemes use the same information. The basis of ordinary adaptive routing, namely the use of average instead of marginal delays, is incorrect regardless of the quality of the available information.

Table II. Message Waiting Times Under Each Routing Scheme.

Source traffic γ	Estimated expected waiting time: Adaptive routing	Estimated expected waiting time: Quadratic routing	% Savings using quadratic routing
A. When message lengths are exponentially distributed			
0.2	1.2435 (0.0410)	1.2239 (0.0348)	1.58
0.5	1.4972 (0.0484)	1.4123 (0.0390)	5.67
1.0	1.9991 (0.0535)	1.8370 (0.0529)	8.11
1.5	3.3176 (0.1528)	3.1136 (0.1631)	6.15
1.8	5.4996 (0.3326)	4.9503 (0.3889)	9.99
B. When message lengths are constant			
0.2	1.1162 (0.00925)	1.1616 (0.00958)	−4.07
0.5	1.3092 (0.0138)	1.3280 (0.0107)	−1.44
1.0	1.6830 (0.0206)	1.5899 (0.0210)	5.59
1.5	2.3919 (0.0297)	2.1445 (0.0431)	10.34
1.8	3.8991 (0.1171)	3.4886 (0.1120)	10.53

6. A Simulation Comparison of Quadratic Routing

Simulations comparing quadratic routing with ordinary adaptive routing have been conducted in order to confirm the model used here and to explore the robustness of the quadratic routing algorithm. Messages were assumed to arrive in a Poisson manner, and message lengths were either exponential (Table IIA) or constant (Table IIB). The simulator used the same network and parameters as the example of Sec. 4, namely $\mu C_1 = \mu C_2 = 1$ and an indirect routing delay of 1.

The estimated waiting times shown in these tables were computed according to the methodology proposed by Fishman [2], which takes advantage of the regenerative property of stable queues to generate sequences of independent, identically distributed observations. The routing simulation took the arrival of a message when both direct and indirect queues were empty as a regeneration point. Estimates and standard errors for the mean waiting time were computed over 1000 regenerations.

The simulation results serve to confirm the superiority of quadratic routing.[2] When message lengths are exponential (as was assumed when developing the model) the savings due to quadratic routing are roughly 5 to 10 percent in magnitude and are greatest when the system is moderately loaded. This is about the same improvement predicted by the theory and it occurs at about the same loads.

Similar savings at moderate to heavy loads are also observed for the polar case of constant message lengths. When γ was small, quadratic routing did not perform as well as ordinary adaptive routing, because it did not correct for the effect of the constant message length. But although the quadratic algorithm was developed using a different assumption about message lengths, it is still useful here.[3]

Note that the comparisons show savings similar to those predicted by the theory despite the fact that the waiting times estimated during the simulations are significantly less than those computed in Table I. This difference is due to the dependence of the arrival rates in each queue on its length relative to the length of the other queue. Either ordinary adaptive routing or quadratic routing will avoid a queue whose length is excessive, thereby reducing the arrival rate when any one queue is long. Just as in Morse's analysis of "customer impatience" [10] this dependence of arrival rate on current load gives a smaller expected value and a more concentrated distribution for the waiting time than the simple model predicts.

7. Summary and Conclusions

This paper has built two simple models of routing in a store-and-forward communications network in order to contrast the optimum routing with the performance of an adaptive routing scheme. The models have shown that adaptive routing, which assigns messages to output channels so as to minimize the apparent message delay, induces equilibrium (stationary) behavior in the network which does not minimize the average message delay. Situations such as this, in which the aggregate of many local optimizations does not yield the global optimum, arise because adaptive routing does not take into account the effect of its current decision on the future state of the network.

The way in which adaptive routing misbehaves has been illuminated by an example which shows that adaptive routing tends to overuse the direct route relative to its alternate. Adaptive routing also initiates use of the alternate route when the traffic was heavier than optimal. The difference in average message delay is greatest at moderate traffic levels, while at light or heavy loads the difference is not so marked.

The form of the optimality condition (8) suggests a modification to the adaptive routing algorithm which would cause it to attain an equilibrium which also satisfied the necessary optimality conditions. The modification computes a quadratic function which estimates the marginal delay per message instead of the average delay. and maintains this estimate in the routing table. This modification requires only a small increase in computational overhead and should cause no increase in line overhead and require no other program changes.

The improvement in performance of the quadratic routing algorithm has been confirmed by simulation. These simulations indicate that the simple model presented here nevertheless was adequate for our analysis of adaptive alternate routing algorithms. Further work, however, remains to be done in relating the predictions of the models and the simulations to the performance of actual systems. Theoretical analysis of more complicated cases involving several destinations, priority messages, and dynamic behavior also needs to be done, and should lead to a better understanding of network routing. From this interaction of theory and practice, routing strategies still better than the one proposed here will doubtless emerge.

Received March 1974; revised February 1975

[2] It has been impossible to test statistically the significance of the differences because the simulation program has not allowed estimation of the degree of correlation between the sample waiting times of different runs. These times are almost certainly positively correlated, not least because the same seed was used for the random numbers generated in each run. The existence of a positive correlation between the waiting times means that the standard error of the difference is substantially less than either of the observed standard errors. Hence the magnitude of the observed differences is more significant than a comparison with these tabulated standard errors would indicate.

[3] Of course a still better algorithm could doubtless be developed from a model which assumed constant message lengths in the first place. Marchand [8] has given the necessary conditions which apply to this case and showed that optimality requires a quadratic bias term, but did not solve for the coefficients.

References

1. Agnew, C.E. The dynamic control of congestion-prone systems through pricing. Stanford U. Ph.D. dis.; Rep. No. 6 of the Program in Information Technology and Telecommunications, Stanford U., Nov. 1973.
2. Fishman, G.S. Statistical analysis for queuing simulations. *Management Sci., 20*, 3 (Nov. 1973), 363–369.
3. Frank, H., and Chou, W. Routing in computer networks. *Networks 1* (1971), 99–112.
4. Fratta, L., Gerla, M., and Kleinrock, L. The flow deviation method; an approach to store-and-forward communication network design. *Networks 3* (1973), 97–133.
5. Fultz, G.L. Adaptive routing techniques for message-switching computer-communication networks. Rep. No. UCLA-Eng-7252 (ARPA), July 1972, UCLA Computer Sci. Dep. Available from NITS as document AD 749 678.
6. Kleinrock, L. *Communication Nets: Stochastic Message Flow and Delay*. Dover, New York, 1972.
7. Luenberger, D.G. *Introduction to Linear and Nonlinear Programming*. Addison-Wesley, Reading, Mass., 1973.
8. Marchand, M. Priority pricing. *Management Sci. 20* (1974), 1131–1140.
9. Meister, B., Mueller, H., and Rudin, H. New optimization criteria for message-switching networks. *IEEE Trans. on Commun. Technol., COM-19* (June 1971), 256–210.
10. Morse, P.M. *Queues, Inventories and Maintenance*. Wiley, New York, 1958.
11. Schwartz, M., and Cheung, C.K. Alternate routing in computer-communication networks. Proc. 7th Hawaii Int. Conf. on System Sci., Jan. 8–10, 1974, U. of Hawaii, pp. 67–69.