Distributed Dynamic Routing Schemes

Peter B. Key
Graham A. Cope

DYNAMIC ROUTING IN A CIRCUIT-SWITCHED network is a way of providing flexibility at the switched level. Flexibility is required to adapt to changing and volatile traffic demands; to cope with forecasting uncertainties, shifts in traffic patterns, and the introduction of new services; and to provide resilience against individual network failures. In other words, dynamic routing, by allowing the path a call takes through the network to vary with time or the state of the network, enables the network to be used more efficiently, and also provides robustness. With performance measured by "grade-of-service"—the probability that a call encounters network congestion—efficiency can translate into a cheaper network for given performance criteria, or into increased performance at a given cost.

A macroscopic example of efficiency gains is given by noncoincident busy hours—different parts of the network can be busy at different times of the day. A good example of this is within the international network, where some countries are asleep while others are at work. This makes the sleeping country amenable to overflow traffic from working countries, and it has been estimated that capacity savings of almost 30% can be achieved by exploiting this noncoincidence with dynamic routing schemes [1].

When a demand is made upon the network, there are essentially two questions to answer. First, should the demand be accepted? Second, where should the demand be routed? In any particular dynamic routing scheme, these two concepts can be inextricably interwoven; however, for convenience we shall say that the first addresses the control issue, whereas the second is concerned with routing. The first question is in some sense more fundamental and harder; it is important because, in the absence of control, the performance of a circuit-switched network can seriously degrade under a flexible routing scheme, which is discussed in the next section.

Dynamic routing is about finding the "best" path through the network, i.e., a short path that causes the least damage to future calls. In its simplest form, dynamic routing is about finding spare capacity in the network. This can be done in a centralized or decentralized way, the former making use of a central processor to monitor the state of the network and recommend decisions. A centralized processor requires a separate control network, and the system is vulnerable to failures. A decentralized or distributed scheme is one in which the intelligence is spread throughout the network. In this article, we concentrate on distributed schemes.

The section following that is a survey of schemes that do not explicitly use much information about the state of the network. In particular, we concentrate on Dynamic Alternative Routing (DAR), which is a simple but highly effective routing method currently planned for the British Telecom Network. In contrast, the next section looks at state-dependent routing, and discusses how some of the methodology also has a bearing on the control issue.

We then briefly touch on the problem of dimensioning a network that uses dynamic routing; in other words, how much capacity do we need, and where should it be put to provide an acceptable performance? A practical example, referring to routing in an international access network, is discussed; and finally, some conclusions are drawn.

Control and Trunk Reservation

A flexible or dynamic routing scheme gives an individual call a better chance of success, perhaps by increasing the number of ways it can traverse the network. However, the success of this call could potentially prevent the success of more than one other call, creating a less than satisfactory solution overall [2]. This is a symptom of the classic dichotomy between individual and social optimization. From the individual call's point of view, we want as much freedom as possible so that it has the best chance of finding a path through the network; whereas to achieve a global optimum (where the carried traffic is at a maximum), we might have to limit an individual's freedom.

For example, consider a fully connected symmetric network such as that in Figure 1, where calls are offered between each node pair and the first-choice route is the direct link joining the two nodes. Now consider any dynamic routing scheme that seeks out a two-link alternative path if the direct route is busy, and suppose that over a long period of time all such two-link al-

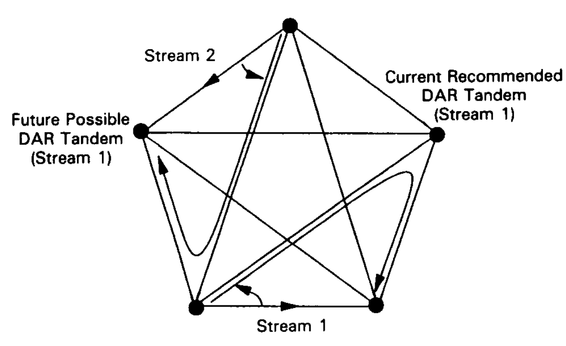

Fig. 1. DAR operation.

0163-6804/90/0010-0054 $01.00 © 1990 IEEE

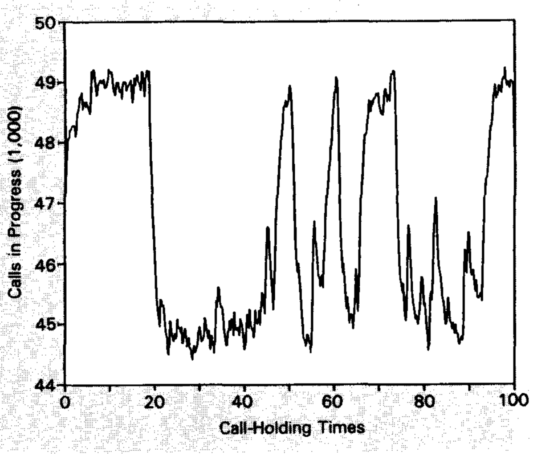

Fig. 2. Instability in an uncontrolled network.

ternatives are chosen an equal number of times. For instance, we could choose a two-link alternative at random, or one of the dynamic schemes such as DAR, described below. Then under-overload calls are set up on two-link paths, which then cause other calls to overflow, boot-strapping the network into a state of high congestion. The network as a whole would be better off restricting calls to one-link paths.

At certain critical traffic levels the network with no controls can exhibit unstable behavior in which it flips between two quasi-stable states. These are high and low blocking states, corresponding to many calls carried on two links and one link, respectively. This behavior was first predicted by simple fixed-point analytic models [3], and has since been found in simulations [4] [5]. Figure 2 shows an example of this behavior, where the total number of calls in progress is plotted for a 5-node fully connected network (10 links) with 5,000 circuits on each link and 4,900 Erlangs (simultaneous calls) of traffic offered between each node pair, where DAR is used to search out a two-link alternative, and where there is no control applied. In the low congestion state, almost 4,900 Erlangs are carried on each route; however, for much of the time the network is in a highly congested state in which each link carries almost 5,000 Erlangs with a significant proportion on two-link paths, and there are only about 4,500 Erlangs carried on each route. This is much worse than if each call were restricted to the one-link paths.

Trunk reservation against two-link overflow calls is an effective way of controlling this, where a trunk reservation of r on a link means that a two-link overflow call is only accepted on that link if there are more than r circuits free, whereas the single-link calls are accepted onto the link whenever there is a free circuit (a trunk reservation of 0). This ensures, first, that under overload two-link alternatives become increasingly unavailable, and second, that performance under any load is always better than if we just used direct routing. Even for large circuit groups of thousands of circuits, only a small trunk reservation parameter is required (less than ten). For instance, in the above example, a trunk reservation parameter of 4 is sufficient to ensure that the network stays in the low congestion state.

Trunk reservation has more general uses, as a way of giving different priorities to different streams of traffic. If a single link is offered to a number of streams of Poisson traffic with a common exponential call-holding time distribution, but which are "worth" different amounts, i.e., generate different revenues if accepted, then applying trunk reservation parameters to each stream is exactly optimal [6] [7]. Here, optimal means that the amount we earn from the link is maximized. If the holding times are not exponential or do not have the same mean, then "approximately" substitutes for "exactly." Similar remarks apply if we are placing constraints on the blocking experienced by the streams, rather than maximizing revenue.

Distributed Dynamic Routing Schemes

Dynamic routing algorithms vary considerably. They may be state-dependent or not, time-dependent or static, and fast or slow (possibly changing decisions for every call arrival or periodically). For example, AT&T's Dynamic Nonhierarchical Routing (DNHR) is a centralized time-dependent scheme, making use of different time zones to prescribe a set of fixed alternate routes for different hours [8]. In general, the faster the response of the algorithm the better; however, there are obvious computational and complexity issues associated with increasingly fast updates.

Distributed schemes can be further differentiated by the amount of information they require. For instance, there are the two extremes of almost "blind" schemes, which make use of negligible information, and complex perspicacious schemes, which require extensive information-passing. Full information-passing is likely to be impractical (it is essentially equivalent to a centralized solution of the problem), and the goal is to find a good routing algorithm that uses only knowledge that is "local."

The following example, depicted in Figure 3, illustrates a routing problem: Suppose that we have three possible routes through the network from a particular source to a specific destination, which we label by $r = 1,2,3$, each of which is offered its own background traffic. A call comes along and we want to send it along one of these three paths—which one should it take? One way to route a stream of such calls is to use "proportional routing," in which we send a certain proportion of calls along route r. The problem, then, is how to update such proportions in light of changing network conditions, and "learning automata" are one way of achieving this.

Learning Automata

An application of automata to the routing problem is given by Narendra and Mars [9]. Learning automata do not measure the state of the network directly, but rather rely on indirect information. They continually offer calls across the routes r according to a probability distribution p_r, which is updated at discrete time stages according to feedback information regarding call completion or rejection. If b_r is the probability that a call is blocked on route r, then a locally optimal strategy would minimize the expected number of lost calls, $\Sigma p_r b_r$. Schemes can reward a route on which a call is successful and punish a route on which a call fails. A favored scheme is the so-called $L_{R-\varepsilon P}$ scheme, in which the decrease in p_r for a failure is small compared with the increase for a success. For a linear scheme, if route i is chosen at time n and the call is successful, then updating is

$$p_i(n+1) = p_i(n) + a\left[1 - p_i(n)\right]$$

$$p_j(n+1) = (1-a)p_j(n) \qquad j \neq i$$

while if the call fails,

$$p_i(n+1) = (1-\varepsilon)p_i(n)$$

$$p_j(n+1) = \frac{\varepsilon}{r-1} + (1-\varepsilon)p_j(n) \qquad j \neq i$$

where a and ε are adjustable parameters, $0 < a < 1, 0 < \varepsilon < 1$ with ε small compared with a, and a is itself usually small, so that the updating is gradual. In a changing environment, under

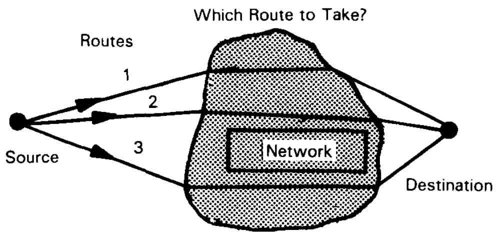

Fig. 3. A routing problem.

certain assumptions [10] [11], it can be shown that the $L_{R-\varepsilon P}$ automata tends to approximately equalize blocking probabilities, b_r, while an L_{R-P} automata, for which $\varepsilon = a$ in the above equalizes the blocking rates $(p_r b_r)$.

Learning automata can also be used to choose the next node rather than a complete path to the destination. This requires a separate updating procedure (automaton) at each node for each destination. The automaton never rejects a call if it is possible to carry it, so a control mechanism can be used to improve performance. Also, no account is taken of the "cost" or length of the path. For example, in a fully connected network it is preferable to first try the direct single-link path joining nodes and then apply the automaton to the choice of two-link overflow paths, rather than use the automaton to include the single-link path.

Implied Cost Methodology

The "implied-cost" methodology of Kelly [12] [13] is another way of calculating proportions for a proportional routing scheme. The rate of return from the network is the sum, over all routes, of the carried traffic on the route multiplied by the "worth" of that traffic. The worths can be arbitrary, but if they are identically one, then the return is just the total carried traffic. Let r be a route index, and k a link index; if link k has capacity C_k, calls worth w_r arrive at Poisson rate λ_r (with unit mean holding time), and the route blocking is L_r, then the rate of return from the network is

$$W(\underline{\lambda}, \underline{C}) = \sum_r w_r \lambda_r (1 - L_r).$$

The approach of Kelly associates an "implied cost" with each link, c_k, which measures the knock-on effects on the entire network of carrying an additional call on that link (in general, there is a separate implied cost for each level of trunk reservation applied to a link). Thus, the true cost of carrying a call on route r is the "surplus value" s_r, the difference between the amount we earn by carrying the call and the damage we do, $s_r = w_r - \Sigma c_k$, where the sum is over all links that r uses. These can be calculated from a set of linear fixed-point equations, and the assumption is that links behave independently, and all traffic is treated as though it were Poisson. For example, in the case of proportional routing where there is no trunk reservation, if b_k denotes the blocking on link k, then under the independence assumption the carried traffic on route r is

$$v_r = \lambda_r (1 - L_r) = \lambda_r \prod_{k \varepsilon r} (1 - B_k)$$

and the traffic effectively offered to link k, ρ_k is given by

$$\rho_k = (1 - B_k)^{-1} \sum_{r: k \varepsilon r} v_r.$$

The fixed-point equations are then

$$c_k = \eta_k (1 - B_k)^{-1} \times \sum_{r: k \varepsilon r} v_r (s_r + c_k)$$

where the sum is over all routes which use link k, with $\eta_k = E(C_k - 1, \rho_k) - E(C_k, \rho_k)$, where C_k is the capacity of link k, and $E(C,a)$ is the Erlangs loss formula for a Erlangs offered to C circuits.

It is possible to calculate the implied costs in a decentralized way in real time—roughly speaking, for each link k we need to know, for each route passing through it, the carried traffic and surplus values. Similarly, we need an associated route calculation where each route knows the implied costs of the links it uses [12]. This can be seen by rewriting the implied cost equations as

$$c_k = \text{const} \times \sum_r \frac{\text{carried traffic on route } r}{\text{carried traffic on link } k} (s_r + c_k)$$

where the constant $\rho_k \eta_k$ is the expected occupancy of the final circuit on link k and can be estimated from the carried traffic.

With the implied costs so calculated, a call offered to route r will generate a net expected revenue of $(1 - L_r)s_r$; thus, traffic should be shared out along routes to reflect these quantities. Therefore, more traffic should be offered to routes with a high net expected revenue—it is possible for $(1 - L_r)s_r$ to be negative, in which case route r should not be used. The change should be gradual, since offering more traffic to a route increases the route blocking and alters the implied costs. However, if we do this, then we have the basis of a decentralized adaptive routing scheme that attempts to maximize the earned revenue (or carried traffic).

DAR

DAR is a simple decentralized control scheme that only uses local information [14–16] and is a natural extension of the currently used Automatic Alternative Routing (AAR), where a number of route choices are attempted sequentially. It was originally designed for application to the British Telecom trunk network, which comprises some 53 Digital Main Switch Units (DMSUs), fully interconnected, shown schematically in Figure 1, with a trunk reservation parameter assigned to each link. A call arriving at a DMSU destined for another DMSU first tries the direct route linking the two exchanges; if this is busy, then the DAR current-choice DMSU is selected and the call attempts to complete via this tandem, with trunk reservation applied against it. If the call is successful, then that tandem stays as the recommended current-choice tandem for the particular source-destination pair, and the next call finding the direct route blocked will also try to complete via this recommended tandem. Conversely, if the call cannot complete via its recommended tandem, then the call is lost and DAR updates the current choice by selecting another DMSU at random from the list of possible DMSUs for that source-destination pair. This is a simple upgrade from AAR, since the DAR choice in the routing table is updated solely on the basis of an appropriate network congestion message.

In the British Telecom trunk network, local exchanges are parented on two DMSUs; thus, there are two ways into the network and two ways out. Therefore, once at a DMSU, there are in fact two possible single-link paths that are attempted first, and then DAR is applied. Similarly, when a call reaches a recommended tandem, there are two possible links it can use to leave the network.

Behavior and Modeling of DAR

One of the desirable properties of DAR is its speed of response. DAR locks onto a good path, and once a path or route ceases to become attractive, another is sought out. This can be thought of as a learning scheme where the probabilities of choosing a particular path are 1, 0, or $1/n$, depending on whether the last call was successful (thus retaining the designated current choice) or not, where n is the number of choices. Although we can gain some idea of DAR's behavior by looking at how it operates for a single source-destination pair, we really need to step back and look at the whole network. On this canvas DAR is attempting, in a distributed manner, to pack calls onto the network, choosing two-link alternatives so that they do not interfere with each other.

The simplicity of the DAR scheme enables mathematical models to be constructed of its long-run stationary behavior. This is important, as it then enables networks to be analyzed and dimensioned. Over a long time period, each DAR choice will be changed an equal number of times, and hence an equal number of calls will be lost on each choice. Thus, if p_t denotes the proportion of calls offered to tandem t, and b_t denotes the proportion of calls lost on choice t, then the blocking rate, $p_t b_t$, is equal for all t (as for the L_{R-P} automaton). Hence,

$$\rho_t \propto \frac{1}{b_t}.$$

An approximation is to calculate b_t using Erlang's formula, with each DAR stream assumed to behave as if it arises from a proportional routing scheme, with proportions p_t. Simulation results have validated the accuracy of this model for many or small DAR streams, but it can break down with a large DAR stream overflowing to a number of small routes, in which case a more delicate analysis is required [17]. Since DAR streams will be kept small in practice, the above simple model is very useful.

Extension to DAR

Although intended for a fully connected network, the principle behind DAR can be extended and applied to different network structures. This principle can be described as a "sticky-random" or "back-the-winner" policy—use what worked last time; otherwise, choose a new alternative at random. Thus, for example, applying the principle to the three-route example of Figure 3, we continue to use route i while calls are not being lost on it, and as soon as we lose a call we choose another route from the r possible at random. This is an example where we can use DAR without a fixed first choice. There are certain other straightforward extensions that can be made, some of which are discussed in [16], and others of which are the subject of current research:

- If the DAR stream is large, it may be preferable to split it. This can be achieved by having a number of choices that are cycled through in turn, with each choice in this "cycle set" independently reset.
- The message to reset the current tandem can be sent before a call is lost. This "DAR reset message" can be returned if, for example, a call routed using DAR takes the last circuit on one of the links of the current choice route. Such anticipation tries to prevent the next call from being lost.
- Certain tandem nodes may be known to be better or worse than others. This can be reflected by increasing or decreasing their probability of selection. For example, some tandems might be barred as a network management action.

Distributed Control and State-Dependent Routing

The simplest form of state-dependent routing, as proposed by Bell-Northern Research [18], routes traffic through links that are "least busy" in the sense of having the largest number of free circuits. More sophisticated schemes attempt to measure the cost of accepting a call on a route. For an isolated group of C circuits offered random (Poisson) traffic of λ Erlangs, this is straightforward. If we assume the calls are "worth" 1 (generate unit revenue), then adding an extra, marked, call when there are j calls in progress costs

$$v(j) = \frac{E(C,\lambda)}{E(j,\lambda)}, \tag{1}$$

where E is Erlang's loss function. This "relative value" $v(j)$ is the expected number of extra calls that are lost when we add a call in state j, and can be thought of as a state-dependent implied cost. Thus, if our marked call is worth w, carrying the call generates a net revenue of $w - v(j)$.

To return to a version of our earlier example, if we have three circuit groups of capacity C_r, $r = 1...3$, and offered random traffic of λ_r Erlangs, and if an isolated call comes along, then we should route the call over the route for which $v_r(j_r)$ is smallest. This is an example of an "index policy," because the index v just depends on the state of its own route. However, if presented with a continuous stream of traffic of λ^* to route over the r choices, then the optimal policy of where to route the arrivals is not in general an index policy, but rather depends on knowledge of the entire state of the network. Nevertheless, we conjecture that an index policy is likely to be close to the optimum, and there are reasons for thinking that an index policy becomes optimal as λ^* is very small [7] or very large (since, in this case, the model approximates a "bandit process" [19], for which an index policy is known to be optimal).

The functions v form the basis of the state-dependent routing scheme proposed by Ott and Krishnan [20] [21]: If we index links by k and assume that links behave independently, then we can estimate the traffic offered to each link λ_k, calculate the relative values v for each link, and approximate the cost of routing a call over r, which uses links $1...l$, by

$$v_r(j_1,...,j_1) = v_1(j_1) + \ldots + v_l(j_l) \tag{2}$$

Here, the effective offered traffic to a link has to take account of blocking it encounters before and after the link in question [22] [23]—the ρ_k mentioned earlier. In a real implementation of such a scheme, it is natural to use two different time scales, making decisions on routing every call arrival, but updating the cost functions v, periodically in time.

If accepting the call is worth w, then we should reject the call if $v_r > w$, since the damage caused by accepting the call is more than the call is worth. Thus, we have a built-in control mechanism. Alternatively, we can use the relative value approach to derive control schemes for other dynamic routing strategies, such as DAR: If a single link is offered several streams of Poisson traffic with the same call-holding time but of different worths, then the optimal control strategy is to use a trunk reservation parameter for each stream. In this case, we can calculate the appropriate relative values [7], and this provides a way of setting control parameters for links in a distributed way on line: We estimate the traffic of each type offered to a link, calculate the relative values, and reject in state j on link k for call type t if $v_k(j) > w_t$. This gives a new trunk reservation parameter, which then feeds back into the calculation of new relative values. These new relative values can also be used in a state-dependent routing scheme using Equation 2 [7].

We have not said how to assign link-based worths to calls. A simple heuristic is to assume that a call using n links is worth $1/n$ on each link it passes through. Some care needs to be taken with calls that can overflow. If we are using an isolated link approach to set trunk reservation parameters (rather than global information via Equations 1 and 2), and if a call can overflow to another path on which it is blocked with probability b, then its effective worth on the original path needs to be decreased by the amount we earn if we reject the call, which is $1 - b$ multiplied by the worth on the alternate path (analogous to the "surplus value" on the alternate path defined in [10]). In the context of a fully connected network using DAR, the first order approximation is that each link is offered fresh traffic worth 1 and overflow traffic worth 1/2, which gives a simple way of setting trunk reservation parameters.

A feature of policies such as trunk reservation is their robustness; i.e., the policies are robust against perturbations of the assumptions. For example, we can alter the offered traffics slightly and still have an almost optimal performance. Also, near-optimal performance can be obtained with a range of trunk reservation parameters; therefore, the exact value is not crucial. This, in turn, means that the exact calculation of the link-based worths is also not essential.

Network Dimensioning

The network dimensioning problem is to provide, at minimum cost, enough capacity to ensure that performance is acceptable, where performance could be an overall average blocking criterion or a requirement on the worst grade of service. In one sense, network dimensioning is simpler with a flexible routing scheme since, provided that there is enough capacity, the dynamic routing scheme should be able to make use of it. Exactly how the capacity is spread around the network is not as critical, so simple rules can be used. Nevertheless, it is possible to maximize the benefit of dynamic routing by using a sympathetic dimensioning scheme, that is, one designed for the particular routing scheme that puts capacity in the best places.

If we pick a particular node in the network, then there needs to be enough capacity leaving the node; in other words, we take the total traffic leaving the node and provide enough capacity so that the grade of service of the total traffic offered to the total number of circuits is acceptable. We then have to allocate the total capacity between routes in some way that could be related to the cost of providing circuits on routes. This is the concept of bundle or bulk dimensioning, which has the advantage that forecasting aggregate loads is easier and more accurate than forecasting smaller individual loads.

The implied cost methodology referred to earlier can form the basis of a sophisticated dimensioning algorithm for which we need to know the derivative of the return, W, with respect to link capacity, C_k. This can be termed the "shadow price," and is a simple linear transformation of the implied costs c_k (the two are equal if there is no trunk reservation). The derivative information provided by the implied costs can be used in a hill-climbing algorithm to dimension the network at minimum cost [24]. The best place to put capacity is on links with a high positive implied cost, where we need to weight the implied cost with the real cost of adding in capacity.

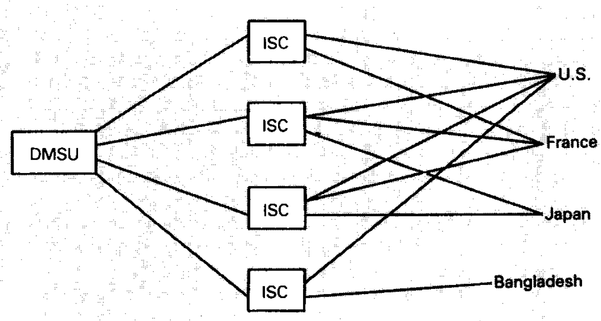

Fig. 4. International access network.

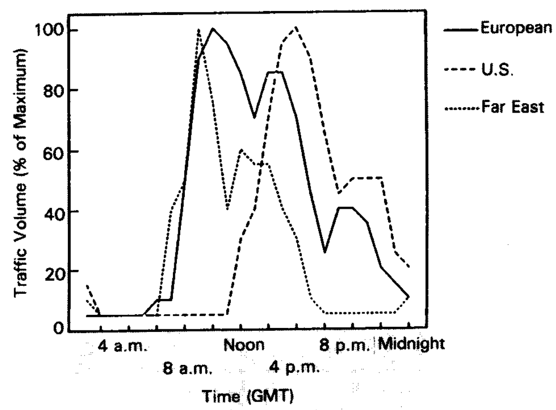

Fig. 5. Outgoing traffic profiles.

A Practical Example

We now discuss the use of dynamic routing schemes in practical networks, looking in particular at the British Telecom international access network.

British Telecom has a number of international gateways, or International Switching Centers (ISCs). These serve to connect the trunk network of DMSUs to foreign administrations. For reliability reasons, those countries with large traffic volumes (for example, the U.S., France, and Germany) are connected to three or four ISCs, as shown in Figure 4. The network between DMSUs and ISCs is termed the "international access network," and every DMSU is connected to every ISC. The routing problem within the access network is to spread traffic over ISCs in such a way that the expensive international circuits are used as efficiently as possible. This is made particularly difficult by the many different traffic profiles the network must handle (see Figure 5), and general uncertainty about exactly which ISC incoming traffic from foreign administrations will be sent to.

There are a number of requirements that the international access network routing scheme must satisfy, including:

- *Ease of implementation*: Network planning overheads and transition from any current scheme must be compensated by the benefits.
- *Robustness*: Once installed, the scheme must be sufficiently flexible to adjust to the varying traffic demands.
- *Efficiency*: Traffic delivered to ISCs by the national network must use expensive international circuits as efficiently as possible. Mismatch between delivered traffic and available circuits should be kept to a minimum.
- *Switching/signaling effort*: Although network administration and circuits are expensive, so are ISCs. Thus, the amount of processing at each switch should be kept to a minimum.

Several routing schemes could be used within such a network, each with their own merits:

- *Least busy alternative*: Each ISC continually signals to every DMSU the number of free circuits it currently has to each destination. Although this may be a good long-term solution, perhaps extended to full state-dependent routing, none of British Telecom's current switches can implement this, and it would have a very large signaling overhead.
- *Proportionate routing*: Each DMSU sends a given proportion of its traffic for each destination country to each ISC. A learning automaton that dynamically adjusts proportions should be very good, but only fixed proportions are possible with British Telecom's current switch technology. Fixed proportions are difficult to set optimally and lack flexibility.

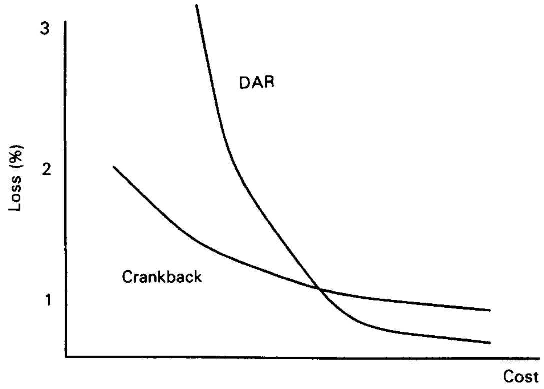

Fig. 6. Afternoon performance during national network and U.S. busy hour.

- *Crankback*: A call is offered in turn to each ISC to which the destination country's ISCs are connected. If a country has circuits on, say, four ISCs, then the switching overhead can be large if a call tries each one. Furthermore, there can be localized areas of congestion, with calls being sent relentlessly to ISCs whose outgoing links are blocked. This inability to balance loads can have significant knock-on effects, as discussed below.
- *DAR with fixed first choice*: The call is offered to a fixed first-choice ISC; but if this is unsuccessful, then a second choice is used, the selection made by DAR. The second choice is invoked if the call fails on the link to the ISC; if crankback is used, then the second choice can be tried if the call also fails on the international link beyond the ISC. A problem here is to allocate the first choices evenly.
- *DAR*: The call is offered to just one current-choice ISC. If the call is successful, then the current choice is retained; otherwise, the call is rejected and the current choice is reset. A disadvantage of this is that each call only has one choice, albeit a good one; but it does balance loads at ISCs very well.

The last option, DAR, satisfies the requirements well, and its ability to choose its own ISCs makes it easy to manage on a day-to-day basis. A particularly desirable feature of DAR is that in all studies to date, adding more capacity into the network has always improved performance. This is not the case with crankback, where the localized congestion problem means that adding more circuits onto a particular national link can increase congestion out of an ISC, causing the performance of some streams (i.e., those which can only use that ISC) to deteriorate. Since the performance measure adopted by British Telecom is the congestion seen by the worst streams, adding this extra capacity could well adversely affect British Telecom's overall performance objective.

Similar effects have been observed when capacity is uniformly added into the access network [25], or for those countries whose busy hour does not coincide with the access network's. Indeed, the two effects are complementary, since adding capacity is equivalent to reducing traffic. The cost/performance curves of Figures 6 and 7 compare crankback and DAR inside the morning (European) and afternoon (U.S. and total network) busy hours, where the performance criteria is the end-to-end congestion seen by the worst 20% of traffic. These graphs highlight the counter-intuitive behavior of crankback.

Current research is investigating hybrids of the aforementioned schemes. One more speculative possibility is to use DAR with fixed first choice, with crankback, and with the first choices determined by proportionate routing—perhaps with the proportions set by a learning automaton.

Dimensioning

Bulk dimensioning is appropriate for deciding how many circuits should be installed in the British Telecom international access network when flexible schemes such as DAR are used. Because the links between a DMSU and ISCs are only used by international traffic to and from that DMSU, there is no noncoincidence of busy hours to exploit, and it becomes sufficient to know only the peak international traffic to and from it. This peak is then split among the ISCs according to, say, the total number of international circuits on each ISC, and then each link is dimensioned using Erlang's formula. The only care that need be taken is that sufficient circuits are installed for traffic that has no flexibility to adjust, such as that which is incoming from foreign countries or outgoing to a country that is only connected to one ISC.

Conclusions

We have described a number of different dynamic routing schemes, all based on decentralized control but varying considerably in flavor. The broad spectrum ranges from the relatively simple DAR, in which routing choices are reset only when a call fails, to more complex state-dependent schemes, which have to know the number of free circuits on a link, or even have to measure or estimate offered traffics. The stage of development of schemes can also be contrasted. Some have been implemented in actual networks or small-scale trials (DNHR), others are being planned (DAR), while some are still undergoing simulation studies (learning automata), whereas the use of implied costs to dynamically control proportions has not even been simulated (to the best knowledge of the authors).

In all cases, the added complexity must be balanced against the cost and ease of administration, monitoring, and performance. Thus, one current research path is focusing on improved cost measures for state-dependent routing. As switches become more powerful, it is possible to implement more sophisticated routing schemes, with the option of performing more complex calculations in the switches or "off-line." But complexity is bought at a price: care needs to be taken that the measures used are robust against modeling assumptions. For example, we wish to be sure that measures for deciding where to route calls based on assumptions of Poisson traffic with the same call-holding time are still appropriate when there is a mix of traffic with different holding times that are not necessarily Poisson.

We also want to know whether more complicated schemes offer worthwhile improvements over simple flexible routing. Simple schemes such as DAR give a surprisingly large benefit and are relatively simple to understand, implement, and con-

Fig. 7. Morning performance during European busy hour.

trol. Moreover, they are not tied to any underlying assumptions about traffic. An interesting open question is the extent to which, as networks grow in size, simple schemes operating quickly and using just local information can provide all the possible benefits. Indeed, there is some evidence that as networks get larger, search techniques that use local random decisions can be highly effective, which is effectively what DAR is doing. Thus, future dynamic routing schemes might become simpler!

References

[1] Y. Watanabe, J. Matsumoto, and H. Mori, "Design and Performance Evaluation of International Networks with Dynamic Routing," *Proc. 11th Int'l. Teletraffic Cong.*, Kyoto, Japan, 1985.

[2] R. Weber, "A Simulation Study of Routing and Control in Communication Networks," *Bell Syst. Tech. J.*, vol. 43, pp. 2,639–2,676, 1964.

[3] Y. Nakagome and H. Mori, "Flexible Routing in the Global Communication Network," *Proc. 7th Int'l. Teletraffic Cong.*, Stolkholm, Sweden, 1973.

[4] J. M. Akinpelu, "The Overload Performance of Engineered Networks with Nonhierarchical Routing," *Proc. 10th Int'l. Teletraffic Cong.*, Montreal, Canada, 1983.

[5] R. G. Ackerley, "Hysteresis-Type Behavior in Networks with Extensive Overflow," *Brit. Telecom Tech. J.*, vol. 5, no. 4, 1987.

[6] S. A. Lippman, "Applying a New Device in the Optimization of Exponential Queuing Systems," *Oper. Res.*, vol. 23, pp. 687–710, 1975.

[7] P. B. Key, "Optimal Control and Trunk Reservation in Loss Networks," *Prob. in the Eng. and Info. Sci.*, vol. 4, pp. 203–242, 1990.

[8] G. T. Ash, R. H. Cardwell, and R. P. Murray, "Design and Optimization of Networks with Dynamic Routing," *Bell Syst. Tech. J.*, vol. 60, no. 8, Oct. 1981.

[9] K. S. Narendra and P. Mars, "The Use of Learning Algorithms in Telephone Traffic Routing—A Methodology," *Automatica*, vol. 19, no. 5, pp. 495–502, 1983.

[10] K. S. Narendra and M. A. L. Thathachar, "On the Behavior of a Learning Automaton in a Changing Environment with Application to Telephone Traffic Routing," *IEEE Trans. Syst. Man and Cybern.*, vol. SMC-20, no. 5, pp. 262–269, May 1980.

[11] P. R. Srikantakumar and K. S. Narendra, "A Learning Model for Routing in Telephone Networks," *SIAM J. Cont. and Optimization*, vol. 20, no. 1, pp. 34–57, May 1980.

[12] F. P. Kelly, "Routing in Circuit-Switched Networks: Optimization, Shadow Prices, and Decentralization," *Adv. Appl. Prob.*, vol. 20, pp. 112–144, 1988.

[13] F. P. Kelly, "Routing and Capacity Allocation in Networks with Trunk Reservation," to appear in *Math. of Oper. Res.*, 1990.

[14] R. R. Stacey and D. J. Songhurst, "Dynamic Alternative Routing in the British Telecom Trunk Network," *ISS '87*, Phoenix, AZ, 1987.

[15] R. J. Gibbens, "Dynamic Routing in Circuit-Switched Networks: The Dynamic Alternative Routing Strategy," Ph.D. thesis, University of Cambridge, England, 1988.

[16] R. J. Gibbens, F. P. Kelly, and P. B. Key, "Dynamic Alternative Routing—Modeling and Behavior," *Teletraffic Science for New Cost-Effective Systems, Networks and Services; Proc. 12th Int'l. Teletraffic Cong.*, Amsterdam: North Holland, 1988.

[17] G. A. Cope, "A Fixed Point Model of Dynamic Alternative Routing," *Brit. Telecom Internal Memo.*, 1988.

[18] W. H. Cameron, J. Regnier, P. Galloy, and A. M. Savoie, "Dynamic Routing for Intercity Telephone Networks," *Proc. 10th Int'l. Teletraffic Cong.*, Montreal, Canada, 1983.

[19] J. C. Gittins, "Bandit Processes and Dynamic Allocation Indices," *J. Royal Stat. Soc.*, vol. B41, pp. 148–164, 1979.

[20] T. J. Ott and K. R. Krishnan, "State-Dependent Routing of Telephone Traffic and the Use of Separable Routing Schemes," *Proc. 11th Teletraffic Cong.*, Kyoto, Japan, 1985.

[21] K. R. Krishnan and T. J. Ott, "Forward-Looking Routing: A New State-Dependent Routing Scheme," *Teletraffic Science for New Cost-Effective Systems, Networks and Services: Proc. 12th Int'l. Teletraffic Cong.*, Amsterdam: North Holland, 1988.

[22] P. M. Lin, B. J. Leon, and C. R. Steward, "Analysis of Circuit-Switched Networks Employing Originating-Office Control with Spill-Forward," *IEEE Trans. on Commun.*, vol. 26, pp. 754–765, 1978.

[23] F. P. Kelly, "Blocking Probabilities in Large Circuit-Switched Networks," *Adv. Appl. Prob.*, vol. 18, pp. 473–505, 1986.

[24] P. B. Key, "Implied Cost Methodology and Software Tools for a Fully Connected Network with DAR and Trunk Reservation," *Brit. Telecom Tech. J.*, vol. 6, no. 3, July 1988.

[25] G. A. Wroe, G. A. Cope, and M. J. Whitehead, "Flexible Routing in the BT International Access Network," *Proc. 7th U. K. Teletraffic Symp.*, IEE, Durham, England, 1990.

Biography

Peter B. Key graduated in mathematics from Oxford University, England, in 1978, and received an M.Sc. in statistics from the University of London, University College, in 1979. He was employed for three years as a Research Assistant at Royal Holloway, working on the relationship between Kalman Filter models and classical forecasting techniques, and was awarded a Ph.D. in mathematical statistics from London University.

He joined the Performance Engineering division of British Telecom's Research Laboratories in 1982, where he has worked on the analysis of circuit-switched networks, and the design and implementation of dynamic routing strategies (in particular, DAR). He currently heads a Network Traffic Management group, which provides consultancy to the Network Management center, performs research in traffic management and control, and is also involved in the design and modeling of present and future transmission networks.

Graham A. Cope graduated in mathematics from Cambridge University, England, in 1986, and received the M.Sc. in information technology from the University of London, Imperial College, in 1987. While an undergraduate at Cambridge, he worked as a Research Assistant under Dr. F. Kelly, with whom he developed efficient computational techniques for the evaluation of implied costs. During his M.Sc. studies, he developed approximate models for multistage interconnection networks as used in parallel computers.

Since then, he has worked at British Telecom's Research Laboratories on the design, control, routing, and planning of telecommunication networks. He has specialized in the modeling of the international network, attempting to find simple (yet close to optimal) planning, measurement, dimensioning, and routing rules, including the best ways to use DAR. Recently, he has been involved in similar studies on British Telecom's national network.