

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
17 January 2002 (17.01.2002)

PCT

(10) International Publication Number  
**WO 02/05495 A1**

(51) International Patent Classification<sup>7</sup>: **H04L 12/56**

(21) International Application Number: PCT/GB01/03012

(22) International Filing Date: 5 July 2001 (05.07.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
00305715.5 6 July 2000 (06.07.2000) EP

(71) Applicant (for all designated States except US): **BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY** [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **TATESON, Jane, Elizabeth** [GB/GB]; 145 High Street, Wickham Market, Woodbridge, Suffolk IP13 0RD (GB). **MARSHALL, Ian, William** [GB/GB]; 6 Pump Close, Bredfield, Suffolk IP13 6DD (GB).

(74) Agent: **ROBINSON, Simon, Benjamin**; BT Group Legal Services, Intellectual Property Department, Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).


(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:  
— with international search report

[Continued on next page]

(54) Title: PACKET ROUTING



(57) Abstract: The present invention provides a routing protocol that scatters a stream of packets along a number of parallel paths, the packets being treated independently. The next hop for each packet is chosen probabilistically by comparing the resistance of available options. The resistance of a given hop depends upon the time the packet would spend in an output buffer from the current node, the time the packet would spend in the input buffer of the next hop node, the transfer time between the nodes and the number of hops that the packet would take from the current node to the ultimate destination of the packet using the shortest path. This routing protocol is more efficient than shortest path first routing under simulation.



WO 02/05495 A1



— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

## PACKET ROUTING

This invention concerns the routing of packets in packet-based networks. The performance of the internet is renowned for its inconsistency. Sometimes a document can be downloaded instantly; at other times, the same document takes a hundred times longer to appear. As a result, there is a lot of research effort directed at reducing internet congestion. Some groups favour resource rationing (K Danielsen and M Weiss, *"User Control Modes and IP Allocation"*, MIT Workshop on Internet Economics, March 1995), other groups seek to increase routing intelligence (G Di Caro and M Dorigo, *"AntNet: Distributed Stigmergetic Control for Communication Networks"*, *J. Artificial Intelligence Research*, **9**, p. 317, 1998) to avoid congested regions of the network. It is generally understood that Shortest Path First (SPF) routing results in unnecessary congestion by focusing all traffic on to the same paths.

The problems associated with SPF are well understood. Some attempts to overcome this problems focus on traffic engineering, for example: Davie *et al.* (*"Optimal use of multiple paths in IP networks"*, IEE 16<sup>th</sup> UK Teletraffic Symposium on 'Management of quality of Service – The New Challenge', Harlow, 22-24 May 2000), Holness and Phillips (*"Dynamic QoS for MPLS Networks"*, IEE 16<sup>th</sup> UK Teletraffic Symposium on 'Management of quality of Service– The New Challenge'), Murphy *et al* (*On Design of Diffserv/MPLS networks to Support VPNs*, IEE 16<sup>th</sup> UK Teletraffic Symposium on 'Management of quality of Service – The New Challenge') and Webb (*"Traffic Engineering in IP Networks: What Does it Offer?"*, IEE 16<sup>th</sup> UK Teletraffic Symposium on 'Management of quality of Service – The New Challenge'). Some of these approaches involve sophisticated schemes for speeding up traffic flow across a network. Currently, the idea of using explicit routes is popular. However, although the use of explicit routes is a workable traffic engineering mechanism, Davie *et al.* do not address how these routes should be chosen initially or under what conditions they should be activated. No criteria for path optimisation are proposed for use with this approach. Optimised paths could be set if traffic flow were predictable, but this is not the case. The Holness and Phillips scheme aims to guarantee QoS by reserving bandwidth for particular classes of traffic. Whilst guaranteed QoS is desirable, reserving

resources is wasteful. Also, their dynamic choice of routes seems to require a heavy overhead of signalling and negotiation. The work by Murphy *et al.* considers traffic load balancing, but does not couple this to convergence on destination.

5           According to a first aspect of the present invention there is provided method of routing a data packet at a network node, the method comprising the steps of:

- (a)       determining the number of valid node outputs that the packet may be routed to;
- 10       (b)       generating a routing factor for each of the valid node outputs that the packet may be routed to; and
- (c)       randomly selecting one of said valid node outputs, wherein the probability of selecting one of the valid node outputs is related to the respective routing factor for that node output. Preferably, the probability of selecting one of the valid
- 15       node outputs is inversely proportional to the respective routing factor for that node output.

The routing factor may depend upon any, some or all of the following: the time that the packet would be buffered in the respective output buffer; the time that the packet would be buffered in the respective input buffer; the packet

20       transmission time to reach the respective node; or the number of hops that the packet would take to reach the ultimate destination of the request using the shortest path.

According to further aspects of the present invention there are provided a network node configured, in use, to operate according to any of the above

25       methods; a communications network comprising one or more of such network nodes; and a data carrier containing computer code for loading into a computer for the performance of any of the above methods.

Advantageously scatter routing gives significantly better throughput than SPF in congested conditions, for a wide range of network connectivities, and

30       performs only slightly less well when the network has no congestion. Another factor that affects relative routing performance is network structure. Whilst there are network topologies that can be devised (*Optimising Network Architectures*, P Bladon, G Chopping, B Jensen and T Maddern, IEE 16<sup>th</sup> UK Teletraffic Symposium

on 'Management of quality of Service – The New Challenge').that are inherently good at balancing traffic load, asymmetric demand for services will still result in under-utilisation of network resources using Shortest Path First. Scatter routing does need sufficient connectivity in order to out-perform SPF. Lack of

5 homogeneity in connectivity *per se* is not a problem, as long as it does not result in connectivity falling so low in parts of the network that parallel paths no longer exist. Simulations show that connectivity with  $m = 25$  is optimal, but connectivity with  $m = 10$  with a variability of 5 is sufficient to give good results. The implication of this could be that nodes in a network with few connections could

10 operate SPF but could pass packets to more richly connected neighbours, where scattering over parallel paths could take place. Connectivity with  $m$  greater than 25 does not improve the performance of scatter routing in our simulation. Even for very large networks, with the simulation parameters as set here, connectivity with  $m = 25$  is sufficient for optimal route parallelisation.

15 Further advantages of the present invention are that scatter routing does not require global knowledge. Needing only local knowledge makes scatter routing dynamically adaptive and reduces the memory needs of the routers, compared with routing protocols that maintain knowledge of multiple end-to-end paths. Scatter routing can also make use of existing SPF routing infrastructure and

20 methodology. The advantage of this is that its implementation would not require radical change to what is already in place. Scatter routing would add an overhead of complexity to the existing routing system. However, compared with some complex schemes that have been proposed, using RSVP and MPLS, our proposal is relatively simple.

25 Scatter routing could be used to differentiate between traffic having different priorities, simply by varying the constant  $k$  associated with a packet. (the greater the value of  $k$  the more closely the route follows the shortest path.)

The invention will now be described, by way of example only, with

30 reference to the following figures in which;

Figure 1 shows a schematic depiction of a communications network which is operated according to the present invention; and

Figures 2-14 shows a schematic depiction of computer telephony terminals connected to a communications network; and

Figure 1 shows an example of the sort of network that can be used with  
5 the method of the present invention. Network 100 comprise a plurality of switching nodes 120 and a plurality of endpoint nodes 110. In the network shown in Figure 1 there are six endpoint nodes 110 and three switching nodes 120, all of which are fully inter-connected (i.e. there is a direct connection for each node to all of the other nodes, whether an endpoint node 110 or a switching  
10 node 120, in the network. As networks grow it becomes infeasible to have a fully interconnected network and each a maximum number of connections,  $m$ , is defined for each network node. The value of  $m$  may be constant for all of the nodes within the network or constrained within a range of values.

In order to demonstrate the advantages of the present invention a number  
15 of different networks were generated in order to perform a number of simulations. The rules used to construct the simulated network were in general as follows. First, 2 switching nodes were positioned randomly in space and linked together, then another switching node was positioned randomly in space and linked to the first two nodes, assuming that  $m \geq 2$ . Connections made from a new node may not  
20 go to the same node more than once. Therefore, until the number of nodes in the network exceeds  $(m + 1)$ , the addition of a new node simply means adding a connection to each of the nodes already in the network. Once the network comprises  $(m + 2)$  nodes, the choice of node to which a new node will connect to is weighted towards those nodes that are already more highly connected. This, to  
25 some extent, simulates the way 'true' networks grow ("*Emergence of Scaling in Random Networks*", A-L Barabasi and R Albert, **Science**, **286**, 1999, 509). In some of the simulations, the choice of which node(s) that a new node would connect to also depended upon its proximity to the newly added node, in order to make the network regions more distinct. The significance of the length of a link,  
30 in the simulation, is that it was used to set the transfer time of packets along the link. The simulations were limited to networks having sizes that varied between 12 and 112 nodes, although it is believed that the present invention would scale to larger sized networks.

At the start of each simulation, the full network capacity is available, that is, there are no requests already in the system. Then, requests arrive at endpoint nodes on the network. In most experiments carried out, the requests arrive at a constant rate, irrespective of the size of the network, which leads to the smallest  
 5 networks becoming more congested than larger networks. The type of request and the endpoint node at which it arrives are chosen randomly.

Having constructed a network, each endpoint node was assigned a number of virtual nodes, *vnode* (the number of which was constant for a given network), each of which was randomly assigned a request handling ability: simply  
 10 one of a possible *maxservices* number of services. Then routing tables were set up at each node for the standard Distance Vector Routing protocol, so that, for each service available on the network, each node knows to which neighbour it should pass any request to ensure the fewest number of hops in order to reach its destination and also the number of hops that this route will take. The routing  
 15 tables were set up to determine a request destination by service, rather than node address.

The premise of the present invention, which will be referred to as scatter routing, is that it is better to use more of the available network resources, rather than simply channelling all requests along the same short routes, as in SPF  
 20 routing. This premise is valid when the network is heavily loaded because congested routes result in requests 'timing out' in queues and overflowing buffers. When the network is uncongested, using the shortest path is optimal, but we argue that the difference between scatter routing and SPF under these conditions is not significant; the critical conditions occur when there is heavy demand, and  
 25 this is when scatter routing is most advantageous.

In the network simulations, a packet flow is divided up into 10 packets which are treated independently (clearly, packet flows can be divided into different number of packets). The next hop for each packet is chosen probabilistically by comparing what we have termed the 'resistance' of available options. The  
 30 calculation of the resistance of a given hop is given as follows:

$$resistance = tbuffer^{\rho_i} + tq_i + ttrans^{\rho_i+k} hop^j_d ,$$

where,

$tbuffer^0_i$  is the time the packet would spend in the output buffer from the current node  $O$ ;

$tq_i$  is the time the packet would spend in the input buffer of the next hop  
5 node  $i$ ;

$ttrans^0_i$  is the transfer time between nodes  $O$  and  $i$ ;

$k$  is a constant; and

$hop^i_d$  is the number of hops that the packet would take from node  $i$  to the destination of the request, using the shortest path.

10

The probability of choosing a given neighbouring node is made inversely proportional to the 'resistance', i.e. the option having the greatest resistance is least likely to be chosen, although this does not mean that it will not be chosen by at least one of the packets in the stream. Thus, when a random number is  
15 generated, the range in which the random number generator can choose a number (for example 1 to 10000) is divided up into sections, one for each of the available routes to a node, each section having a size inversely proportional to the 'resistance' of the route via that node. The section in which the random number falls corresponds to the node to which the packet is passed.

20

Only two constraints are applied to the choice of next hop for a packet. Retracing a step is prevented, unless a dead-end has been reached and an endpoint node that is unable to handle the request will not be chosen. In this way, in general, a neighbouring node is chosen for the next hop if the route to it is uncongested, and if this step makes satisfactory progress towards the packet  
25 destination. In practice, these two constraints require that each node has a knowledge of the services offered by all of the nodes to which it is connected (which is not a serious memory implication), and packets (or a packet at the head of a flow, for larger scattered units) would have to have header space to maintain a list of the nodes through which they had already passed.)

30

The choice of route is made at each node along the way. Because decisions are made without global knowledge and are made probabilistically, rather than absolutely, scatter routing should not result in routes oscillating unstably, as is possible in routing methods that rely on congestion measurement.



Scatter routing implies the overhead of knowing the lengths of queues on connected nodes, and of knowing the number of shortest path hops for each service from each connected node, rather than just from the originating node, as in SPF. This will probably limit the scalability of scatter routing. However, the  
5 scaling would not be a problem within a core network, and if the core network operated efficiently this would have great benefits for the operation of the internet in general.

The corresponding advantage of having routing knowledge for a particular destination via all a node's neighbours, is that the routing operation is not  
10 disrupted by convergence delays in the routing tables after a link failure. With SPF alone, it can take a significant length of time for the knowledge of a failure to propagate through the system. During this time, packets continue to be forwarded along routes towards the failed link, and then time out waiting for a new route to be found. But, with scatter routing, not only will packets utilise multiple paths  
15 (and therefore have less chance of encountering the broken link, for a given destination) but, arriving at a node where the next hop is unexpectedly broken is not disastrous. The routing algorithm will simply route around the broken link, because the use of all possible routes is normal. As the packet is prevented from retracing its steps, it is forced to discover a new route around the obstacle. This  
20 property might greatly reduce the frequency of update messages needed, compared with standard Distance Vector routing, and prove a significant advantage.

Additionally, scatter routing results in a greater proportion of the network being used than with SPF, and, because the load is spread better, there is much  
25 less chance of local congestion arising than when using Shortest Path First.

The 'resistance' function determines the next hop without the need for a heavy overhead of messages or agents to monitor performance. Scatter routing optimises the choice of paths without the need to predict traffic flow. In addition, the scatter routing algorithm responds to traffic flow dynamically, with the speed  
30 of adaptation limited by the rate of update of queue lengths. The paths followed are evaluated node by node, and therefore respond much faster to traffic variation than do paths in schemes that choose routes at the ingress to a network region.

The simulation models splitting up requests of 10 packets into their constituent packets, and making routing decisions packet by packet. This level of control is probably too great for realistic implementation, but the principles tested here would still be valid for dividing up flows into much larger pieces, as long as the relative splitting, compared with SPF, remained the same.

Furthermore, resources do not need to be pre-allocated, so bandwidth capacity is used efficiently for all classes of traffic. If differentiating traffic is important, this could be achieved within scatter routing by labelling a packet with a priority that would determine the constant  $k$  used for forwarding. The greater the value of  $k$ , the more closely the routing follows Shortest Path First. In this way, high priority packets could be routed along the shortest paths, whilst others avoid contributing to local congestion by taking longer routes. Note that, for scatter routing to work, each node needs to know the lengths of input queues on its neighbours accurately, within the time frame of the routing decision. It has been assumed that this would be possible, without testing a method or evaluating its performance. In addition, during this work, we have not been concerned with the ordering of packets within a request, or the re-sending of lost packets. Whereas SPF naturally leads to consistent packet order, scatter routing does not. We have assumed that the arrival window of scattered packets is small enough that the time needed for re-ordering does not impair performance.

The first set of experiments compared request handling throughput for simulations networks of varying size with the same rate of request arrival. Therefore, the smaller the network, the greater the congestion. In these simulations, a request of 10 packets arrived at one of the endpoint nodes every timestep for 200 timesteps, followed by 12 requests arriving every 20 timesteps.

The main focus in giving results has been on comparing throughput during a finite period when there is a surge in demand on the network. The latency of a request is measured as the time taken for all ten packets of a request to be handled at a distant node, and then passed on to a request destination. This request destination is a node, chosen at random, when the request is first generated. These two stages of the request's progress are independent. Note that, in Figure 2, the results presented only include requests for which there has been zero packet loss.

In the first simulation, the results of which are shown in Figure 2), the two lines represent the throughput of networks that are fully connected in a homogeneous way. For example, if the network has 32 endpoint nodes and 30 switching nodes, then  $m$  will be 30. Note that, as the network increases in size, so does the number of services that can be requested. The services are distributed randomly over the virtual nodes on the endpoint nodes. Each endpoint node has a number of virtual nodes such that, on average, 20% of its traffic will be handled at the node where the request originates

The results presented in Figure 2 show that, as the network increases in size, i.e., as its capacity increases, the performance of both scatter routing (the line connecting the crosses) and SPF (represented by the plain line) improves. However, scatter routing attains close to optimal performance very rapidly, whereas SPF continues to result in relatively low throughput, because of a large proportion of requests in which packets are lost.

The plots with square symbols repeat the above simulations with an  $m$  value of 25, i.e. a network that is not fully interconnected. Figure 2 shows that the performance of scatter routing is very similar to that with a fully connected network, whereas SPF seems to show a slight worsening of throughput as the size of the network increases. Again, scatter routing provides significantly increased throughput when compared to SPF routing. This illustrates that fully connected networks are not necessary for scatter routing to perform well. In fact, it is the performance of SPF that suffers more by reducing connectivity, because of the creation of bottlenecks that lead to congestion.

Figure 3 shows the frequency distribution of latencies for a network having 27 endpoint nodes, 25 switching nodes and an  $m$  value of 25 (i.e. a fully interconnected node), with latencies being recorded for each packet. This gives a somewhat different impression from the results shown in Figure 2, because we are not discounting packets that are in a request in which packets have been lost. Figure 2 shows that the throughput for SPF appears much better, but that scatter routing still performs better than SPF. With scatter routing, there is no packet loss, whereas nearly 40% of packets are lost using SPF routing. The latency frequency plot for completed requests, for the same network and simulation conditions, is given in Figure 4.

Another way of illustrating the conditions in which scatter routing performs well is to plot request handling throughput against the interval between request arrivals. This is just another way of varying the level of congestion, again using the same network size and connectivity (i.e. 27 endpoint nodes, 25 switching nodes and an  $m$  value of 25). The results are shown in Figure 5 with each data point being the average of 2 simulations. Figure 5 indicates that the request throughput for scatter routing remains consistently high for all request arrival intervals, whilst the throughput for SPF routing falls off significantly as the request arrival interval decreases. Note that, for all request arrival intervals, there is no packet loss in these simulations for scatter routing, whereas, using SPF, there is a packet loss of 1% for requests arriving at intervals of 10 and 15 timesteps, 2% for requests arriving at intervals of 4, 5 and 7 timesteps, and between 10% and 20% for requests arriving at intervals of 2 or 3 timesteps. When the request arrival interval falls to 1 timestep, SPF routing has a packet loss of nearly 40%. Again, we are seeing that scatter routing enables congestion to be diffused over the network, leading to fewer pressure points or 'hot spots', and consequently better request throughput, in spite of using longer routes.

Scatter routing gives a consistently high throughput, even in very congested conditions. However, when we split a request into a number of packets and scatter it over the network, it may be important to know how large the window in which packets from a given request arrive. Figure 6 shows the average size of window, in which the packets of a request arrive at their final destination. As this only refers to requests in which all its packets arrive, we would expect the SPF window to be smaller. However, we see that, in fairly congested conditions (requests being generated at intervals of 2,3 or 4 timesteps), the window in which packets of a request arrive is not much larger for scatter routing than for SPF. In the most congested conditions, the window for scattered packets is twice the size of SPF packets. However, this result is not surprising, as the only SPF requests that are successful are those that complete their journeys before the network becomes congested.

So far, the simulations have only considered networks that are connected very richly and homogeneously. Figures 7 to 13 display simulation results for networks that have varying levels of connectivity. The simulation results for the

were generated using a network having 27 endpoint nodes and 25 switching nodes, using  $k=30$  and with a request arrival interval of 1 timestep.  $m$  is chosen randomly for each node from a specified range of values, with a different range being used for each simulation. Each data point on the graphs of frequencies of latency was generated as the average of 3 simulations and only complete requests of 10 packets were recorded.

Figure 7 shows a typical comparison of scatter routing versus Shortest Path First during a surge in demand on a network of 27 endpoint nodes, 25 switching nodes and with the connectivity given by  $20 < m \leq 25$  (in order to keep the average load on the nodes approximately the same for these simulations the average request arrival intervals were increased in proportion to the decrease in connectivity). This network is still highly connected with slightly reduced homogeneity when compared with the networks that have been simulated and discussed above. Although SPF handles slightly more requests than scatter routing at low latency, as the surge in demand continues, the SPF handling rate drops while scatter routing continues to handle requests, so that, overall, scatter routing handles far more requests than SPF.

In Figure 8, the network connectivity is chosen to be in the range  $15 < m \leq 20$ , giving a network that is less homogeneous and less highly connected. The slightly longer request arrival intervals allow both SPF and scatter routing to have greater throughput at lower latency, but this effect is most marked for scatter routing, with the window for requests arriving at their final destinations now comparable with SPF. Clearly, the reduction in connectivity does not impair the performance of scatter routing. In Figure 9, the connectivity is given by  $10 < m \leq 15$ . This is a significant reduction in homogeneity and connectivity. However, scatter routing is unaffected by this change, with a very similar throughput to that shown in Figures 7 & 8.. Furthermore, the window in which requests arrive at their final destinations is now greater for SPF than scatter routing.

In Figure 10, the network connectivity range is  $5 < m \leq 10$ . Although scatter routing has to route through some of the same bottlenecks that SPF experiences, it has sufficient alternative routes to give a much higher relative throughput than SPF. The reduced connectivity means that SPF has to route

through more hops to reach its destinations, and this reduces its throughput, as well as increasing the window in which requests arrive at their final destinations, which is now significantly larger for SPF than for scatter routing. In Figure 11, the network connectivity is given by  $2 < m \leq 7$ . Under these conditions, scatter routing  
 5 begins to lose its advantage over SPF, both in terms of the throughput and in the window in which requests arrive at their final destinations, although both are still better when compared with scatter routing. Some packets (about 1%) are lost with scatter routing, compared with approximately 40% with SPF.

In Figure 12, the network connectivity is given by  $1 < m \leq 6$ . This results in  
 10 a significant reduction in throughput with scatter routing, and enlarges the window in which requests arrive at their final destination to 20% greater than with SPF. The packet loss with scatter routing is about 6%, compared with packet loss of nearly 50% for SPF. This connectivity is probably the minimum for scatter routing to operate, in its current form. The relative throughput results for the simulation  
 15 results shown in Figures 7 to 12 are summarised in Figure 13.

Figure 14 shows the result of a simulation of a network of 302 endpoint nodes and 200 switching nodes. In this case, there were 10 'virtual nodes' on each endpoint node (*i.e.* 10 services available), with a total of 50 services for the system as a whole. 512 requests of 10 packets each were generated at intervals  
 20 of 1 timestep initially, rising to intervals of 20 timesteps at time 200 timesteps. Comparison was made between scatter routing and Shortest Path First routing, as in the above discussion, but this time, scatter routing was implemented as described above and also without incorporating the next hop load,  $tq_i$ , in the modified resistance factor, *i.e.*,

25

$$Resistance_{mod} = tbuffer^0_i + ttrans^0_i + k hop^i_d ,$$

where,

$tbuffer^0_i$  is the time the packet would spend in the output buffer from the  
 30 current node  $O$ ;

$ttrans^0_i$  is the transfer time between nodes  $O$  and  $i$ ;

$k$  is a constant; and

$hop^i_d$  is the number of hops that the packet would take from node  $i$  to the destination of the request, using the shortest path.

Each point shown in Figure 14 is the average of 3 simulations. The  
5 scatter routing implemented using the standard resistance factor is represented by  
the line connecting the crosses, the scatter routing implemented using the  
modified resistance factor is represented by the triangles and the SPF routing is  
represented by the line connecting the circles. Figure 14 shows again the much  
greater throughput that can be achieved by using scatter routing. We also see  
10 that the performance of scatter routing is not significantly affected by using  
 $resistance_{mod}$  when determining the routing of a packet as opposed to  $resistance$ .  
Thus, for these simulated networks, knowledge of the next hop input queue load  
is not necessary for the success of scatter routing. Additionally, it should be  
noted that scatter routing in this simulation has zero packet loss, whereas SPF  
15 leads to about 50% packet loss and that the window in which packets arrive at  
their final destinations is slightly smaller for scatter routing than for SPF.

## CLAIMS

1. A method of routing a data packet at a network node, the method comprising the steps of:
  - 5 (a) determining the number of valid node outputs that the packet may be routed to;
  - (b) generating a routing factor for each of the valid node outputs that the packet may be routed to;
  - (c) randomly selecting one of said valid node outputs, wherein the  
10 probability of selecting one of the valid node outputs is related to the respective routing factor for that node output.
  
2. A method of routing a data packet at a network node according to claim 1, wherein in step (c) the probability of selecting one of the valid node outputs is  
15 inversely proportional to the respective routing factor for that node output.
  
3. A method of routing a data packet at a network node according to claim 1 or claim 2, wherein the routing factor depends upon the time that the packet  
20 would be buffered in the respective output buffer.
  
4. A method of routing a data packet at a network node according to any preceding claim, wherein the routing factor depends upon the time that the packet would be buffered in the respective input buffer.
  
- 25 5. A method of routing a data packet at a network node according to any preceding claim, wherein the routing factor depends upon the packet transmission time to reach the respective node.
  
6. A method of routing a data packet at a network node according to any  
30 preceding claim, wherein the routing factor depends upon the number of hops that the packet would take to reach the ultimate destination of the request using the shortest path.



7. A network node configured, in use, to operate according to the method of any of claims 1 to 6.
8. A communications network comprising one or more network nodes  
5 according to claim 7.
9. A data carrier containing computer code for loading into a computer for the performance of the method of any of claims 1 to 6.




Figure 1




Figure 2




Figure 3




Figure 4




Figure 5

Fig 1.5 Window in which request is completed, for a given request arrival interval, on a network of 52 nodes, 25 switches,  $m=25$ ,  $k=30$




Figure 6




Figure 7






Figure 8




Figure 9




Fig 2.4 Latency frequency; 52 nodes, 25 switches; 5 < m <= 10.

Figure 10




Figure 11




Figure 12

Fig 2.7 Ratios of throughput (scatter/SPF) for  $(m-5) < \text{connectivity} \leq m$ .




Figure 13




Figure 14

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 01/03012

A. CLASSIFICATION OF SUBJECT MATTER  
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 974 224 A (BOONE STEPHEN W) 27 November 1990 (1990-11-27)	1-3,5, 7-9
Y	column 2, line 21 - line 50	6
A	column 9, line 6 - line 26 ---	4
Y	EP 0 858 189 A (HITACHI LTD) 12 August 1998 (1998-08-12)	6
A	WO 00 24164 A (AMADON & ASSOCIATES INC ;SOWIZRAL HENRY ADAM (US); ZIKAN KAREL (US) 27 April 2000 (2000-04-27) page 3, line 1 - line 9 page 5, line 17 -page 6, line 15 page 8, line 21 -page 10, line 11 --- -/--	1,2,4

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents :

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

- \*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \* & \* document member of the same patent family

Date of the actual completion of the international search

7 November 2001

Date of mailing of the international search report

14/11/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Lai, C



## INTERNATIONAL SEARCH REPORT

Int ional Application No

PCT/GB 01/03012

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>SCHOONDERWOERD R ET AL: "ANT-BASED LOAD BALANCING IN TELECOMMUNICATIONS NETWORKS" ADAPTIVE BEHAVIOR,US,MIT PRESS, CAMBRIDGE, MA, US, vol. 5, no. 2, 21 September 1996 (1996-09-21), pages 169-207, XP000656299 ISSN: 1059-7123 figure 4</p> <p>-----</p>	1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 01/03012

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 4974224	A	27-11-1990	NONE		
EP 0858189	A	12-08-1998	JP 10224400 A		21-08-1998
			EP 0858189 A2		12-08-1998
			US 6112248 A		29-08-2000
WO 0024164	A	27-04-2000	US 6310881 B1		30-10-2001
			AU 1207300 A		08-05-2000
			WO 0024164 A1		27-04-2000