




(43) **Pub. Date:** **Feb. 25, 2010**






FIG. 1




FIG. 2




FIG. 3

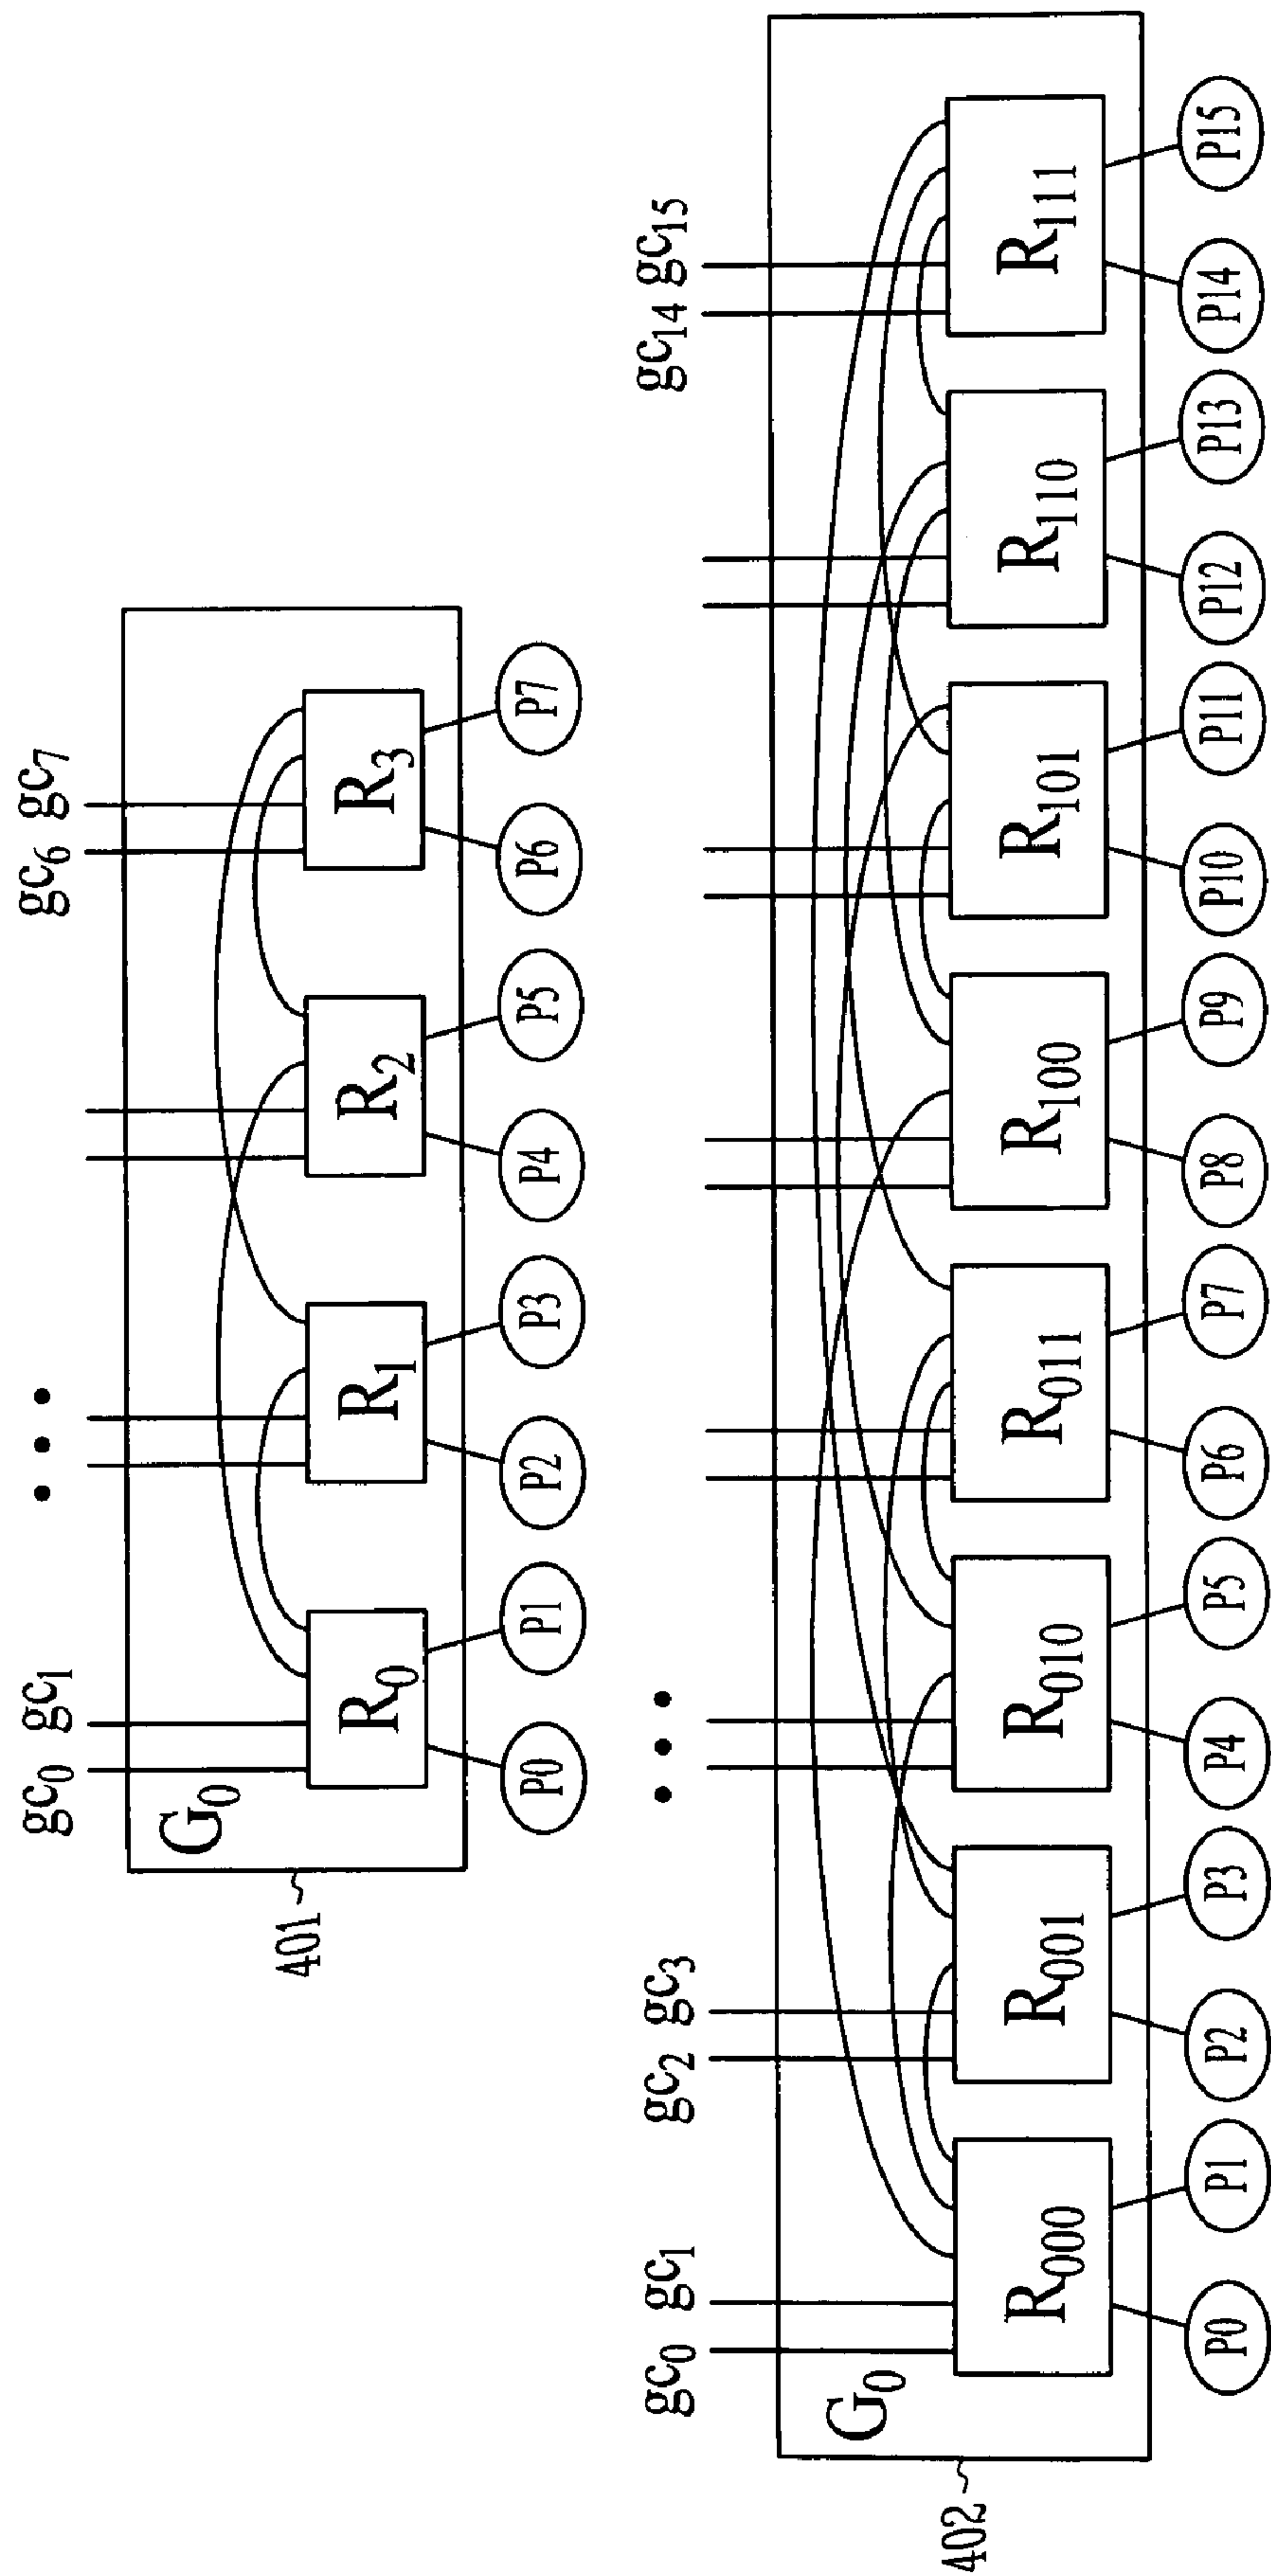


FIG. 4




FIG. 5




FIG. 6




FIG. 7

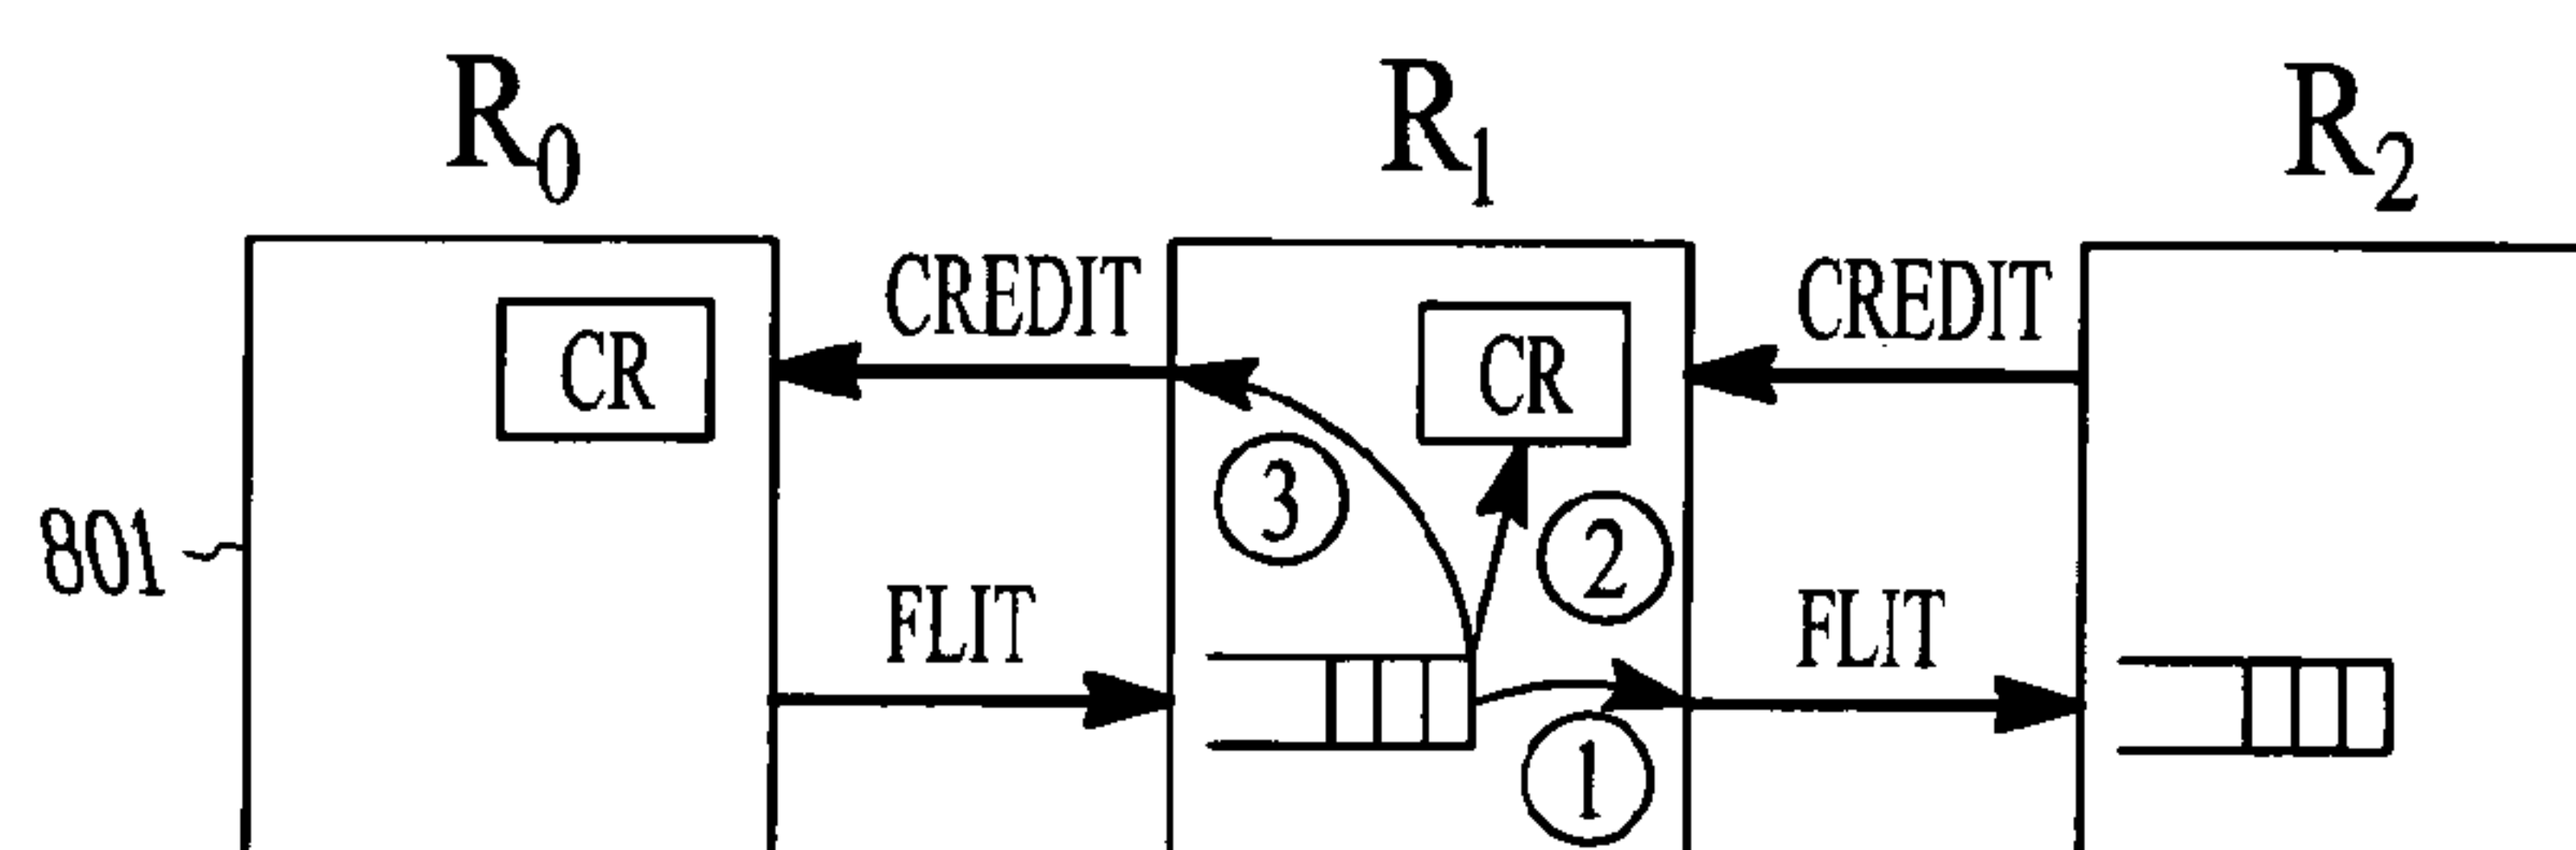


FIG. 8A

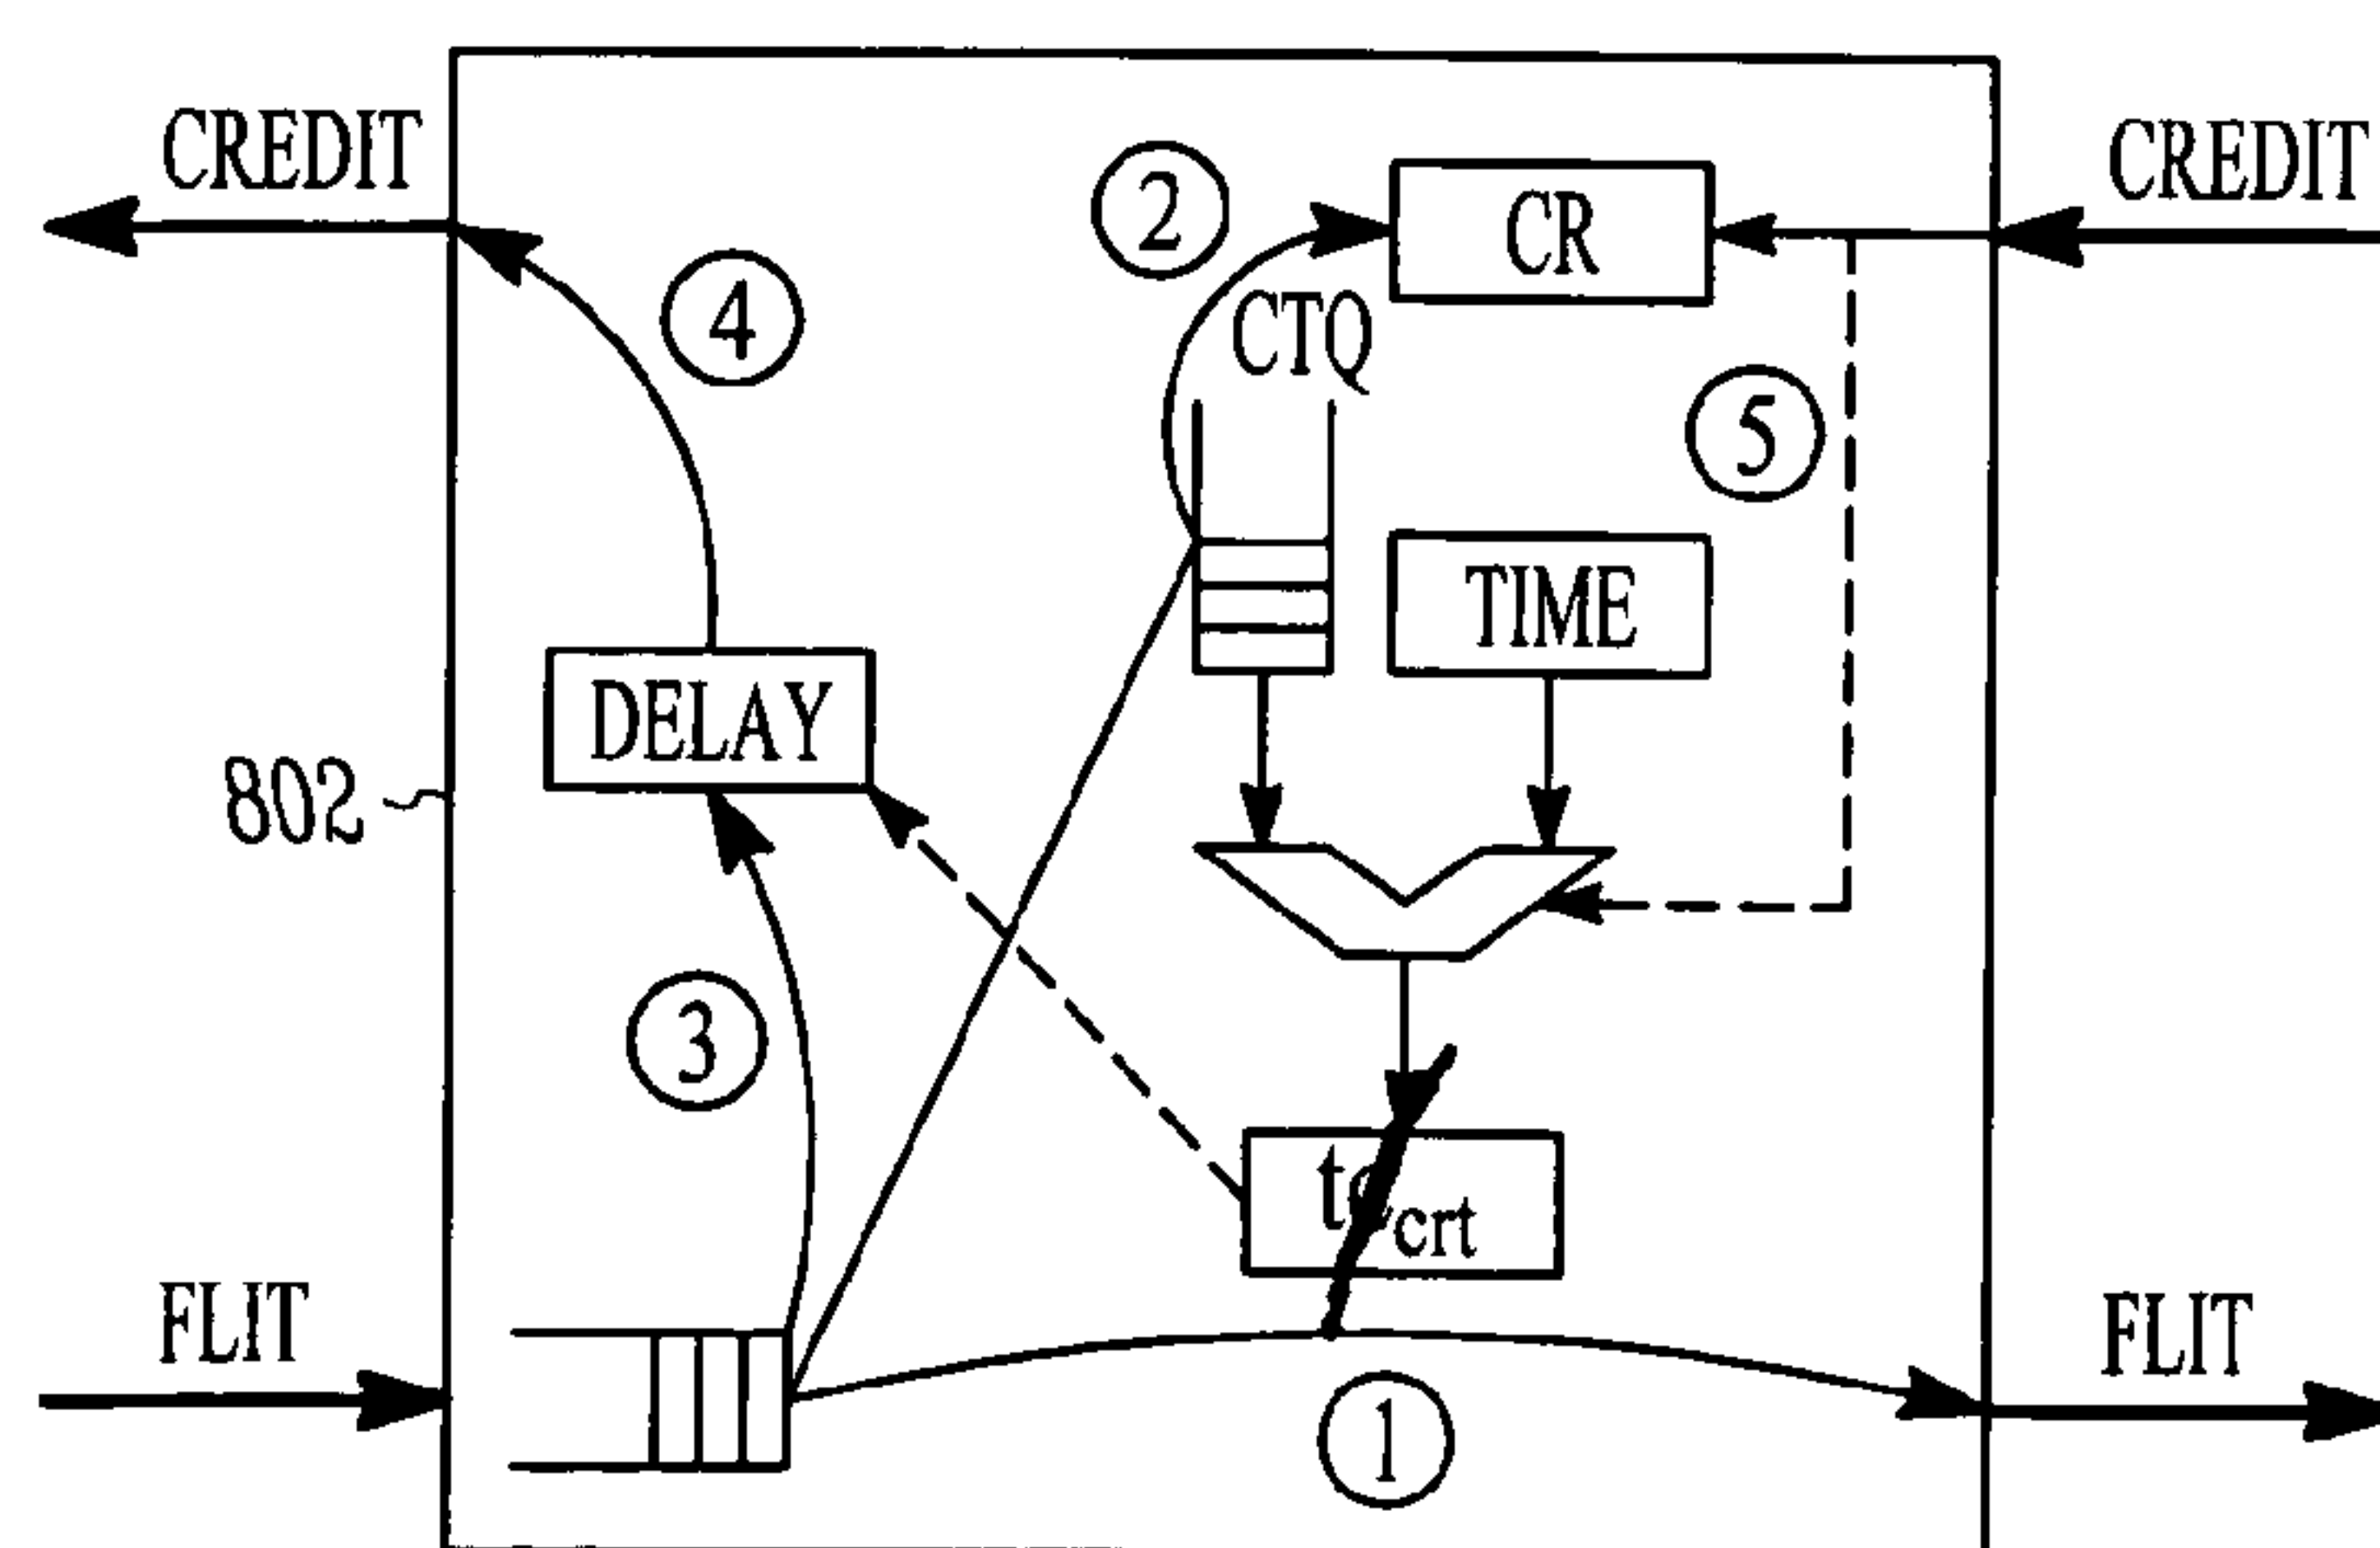


FIG. 8B

DRAGONFLY PROCESSOR INTERCONNECT NETWORK

FIELD OF THE INVENTION

[0001] The invention relates generally to computer interconnect networks, and more specifically in one embodiment to a dragonfly topology processor interconnect network.

LIMITED COPYRIGHT WAIVER

[0002] A portion of the disclosure of this patent document contains material to which the claim of copyright protection is made. The copyright owner has no objection to the facsimile reproduction by any person of the patent document or the patent disclosure, as it appears in the U.S. Patent and Trademark Office file or records, but reserves all other rights whatsoever.

BACKGROUND

[0003] Computer systems have long relied on network connections to transfer data, whether from one computer system to another computer system, one computer component to another computer component, or from one processor to another processor in the same computer. Most computer networks link multiple computerized elements to one another, and include various functions such as verification that a message sent over the network arrived at the intended recipient, confirmation of the integrity of the message, and a method of routing a message to the intended recipient on the network.

[0004] Processor interconnect networks are used in multiprocessor computer systems to transfer data from one processor to another, or from one group of processors to another group. The number of interconnection links can be very large with computer systems having hundreds or thousands of processors, and system performance can vary significantly based on the efficiency of the processor interconnect network. The number of connections, number of intermediate nodes between a sending and receiving processing node, and the speed or type of connection all play a factor in the interconnect network performance.

[0005] Similarly, the network topology, or pattern of connections used to tie processing nodes together affects performance, and remains an area of active research. It is impractical to directly link each node to each other node in systems having many tens of processors, and all but impossible as the number of processors reaches the thousands.

[0006] Further, the cost of communications interfaces, cables, and other factors can add significantly to the cost of poorly designed or inefficient processor interconnect networks, especially where long connections or high-speed fiber optic links are required. A processor interconnect network designer is thereby challenged to provide fast and efficient communication between the various processing nodes, while controlling the number of overall links, and the cost and complexity of the processor interconnect network.

[0007] The topology of a network, or the method used to determine how to link a processing node to other nodes in a multiprocessor computer system, is therefore an area of interest.

SUMMARY

[0008] The invention comprises in one example a dragonfly topology network, comprising a plurality of processor nodes, a plurality of routers, each router directly coupled to a plural-

ity of terminal nodes, the routers coupled to one another and arranged into a group, and a plurality of groups of routers, such that each group is connected to each other group via at least one direct connection.

[0009] Network data is routed in some embodiments using at least one of credit round-trip latency as an indicator of channel congestion and selective virtual channel discrimination.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a block diagram of a dragonfly network topology, consistent with an example embodiment of the invention.

[0011] FIG. 2 is a graph illustrating scalability of a dragonfly network in nodes for various router radices, consistent with an example embodiment of the invention.

[0012] FIG. 3 is a block diagram illustrating a dragonfly network topology, consistent with an example embodiment of the invention.

[0013] FIG. 4 is block diagram of dragonfly network topology groups, consistent with some example embodiments of the invention.

[0014] FIG. 5 is a block diagram of a dragonfly network illustrating minimal and non-minimal routing using virtual channels, consistent with an example embodiment of the invention.

[0015] FIG. 6 is a graph illustrating latency v. offered load for a variety of routing algorithms using various traffic patterns, consistent with an example embodiment of the invention.

[0016] FIG. 7 is a node group diagram of a dragonfly topology network illustrating adaptive routing via global channels using backpressure from intermediate nodes, consistent with an example embodiment of the invention.

[0017] FIG. 8 is a node diagram illustrating credit round trip latency tracking, consistent with an example embodiment of the invention.

DETAILED DESCRIPTION

[0018] In the following detailed description of example embodiments of the invention, reference is made to specific examples by way of drawings and illustrations. These examples are described in sufficient detail to enable those skilled in the art to practice the invention, and serve to illustrate how the invention may be applied to various purposes or embodiments. Other embodiments of the invention exist and are within the scope of the invention, and logical, mechanical, electrical, and other changes may be made without departing from the subject or scope of the present invention. Features or limitations of various embodiments of the invention described herein, however essential to the example embodiments in which they are incorporated, do not limit the invention as a whole, and any reference to the invention, its elements, operation, and application do not limit the invention as a whole but serve only to define these example embodiments. The following detailed description does not, therefore, limit the scope of the invention, which is defined only by the appended claims.

[0019] Interconnection networks are widely used to connect processors and memories in multiprocessors, as switching fabrics for high-end routers and switches, and for connecting I/O devices. As processor and memory performance continues to increase in a multiprocessor computer system,

the performance of the interconnection network plays a central role in determining the overall performance of the system. The latency and bandwidth of the network largely establish the remote memory access latency and bandwidth.

[0020] A good interconnection network typically designed around the capabilities and constraints of available technology. Increasing router pin bandwidth, for example, has motivated the use of high-radix routers in which increased bandwidth is used to increase the number of ports per router, rather than maintaining a small number of ports and increasing the bandwidth per port. The Cray Black Widow system, one of the first systems to employ a high-radix network, uses a variant of the folded-Clos topology and radix-64 routers—a significant departure from previous low-radix 3-D torus networks. Recently, the advent of economical optical signaling enables topologies with long channels. However, these long optical channels remain significantly more expensive than short electrical channels. In this paper, we introduce a Dragonfly topology, that exploits emerging optical signaling technology by grouping routers to further increase the effective radix of the network.

[0021] The topology of an interconnection network largely determines both the performance and the cost of the network. Network cost is dominated by the cost of channels, and in particular the cost of the long, global, inter-cabinet channels. Thus, reducing the number of global channels can significantly reduce the cost of the network. To reduce global channels without reducing performance, the number of global channels traversed by the average packet must be reduced. The dragonfly topology introduced in this paper reduces the number of global channels traversed per packet using minimal routing to one.

[0022] To achieve this global diameter of one, very high-radix routers, with a radix of approximately $2\sqrt{N}$ (where N is the size of the network) are used. While radix 64 routers have been introduced, and a radix of 128 is feasible, much higher radices in the hundreds or thousands are needed to build machines that scale to 8K-1M nodes if each packet is limited to only one global hop using traditional very high radix router technology. To achieve the benefits of a very high radix with routers without requiring hundreds or thousands of ports per node, the Dragonfly network topology proposes using a group of routers connected into a subnetwork as one very high radix virtual router. This very high effective radix in turn allows us to build a network in which all minimal routes traverse at most one global channel. It also increases the physical length of the global channels, exploiting the capabilities of emerging optical signaling technology.

[0023] Achieving good performance on a wide range of traffic patterns on a dragonfly topology involves selecting a routing algorithm that can effectively balance load across the global channels. Global adaptive routing (UGAL) can perform such load balancing if the load of the global channels is available at the source router, where the routing decision is made. With the Dragonfly topology, however, the source router is most often not connected to the global channel in question. Hence, the adaptive routing decision is made based on remote or indirect information.

[0024] The indirect nature of this decision leads to degradation in both latency and throughput when conventional UGAL (which uses local queue occupancy to make routing decisions) is used. We propose two modifications to the UGAL routing algorithm for the Dragonfly network topology that overcome this limitation with performance results

approaching an ideal implementation using global information. Adding selective virtual-channel discrimination to UGAL (UGAL-VC H) eliminates bandwidth degradation due to local channel sharing between minimal and non-minimal paths. Using credit-round trip latency to both sense global channel congestion and to propagate this congestion information upstream (UGAL-CR) eliminates latency degradation by providing much stiffer backpressure than is possible using only queue occupancy for congestion sensing.

[0025] High-radix networks reduce the diameter of the network but require longer cables compared to low-radix networks. Advances in signaling technology and the recent development of active optical cables facilitate implementation of high-radix topologies with longer cables.

[0026] An interconnection network is embedded in a packaging hierarchy. At the lowest level, the routers are connected via circuit boards, which are then connected via a backplane or midplane. One or more backplanes are packaged in a cabinet, with multiple cabinets connected by electrical or optical cables to form a complete system. The global (inter-cabinet) cables and their associated transceivers often dominate the cost of a network. To minimize the network cost, the topology should be matched to the characteristics of the available interconnect technologies, such as cost and performance.

[0027] The maximum bandwidth of an electrical cable drops with increasing cable length because signal attenuation due to skin effect and dielectric absorption increases linearly with distance. For typical high-performance signaling rates (10-20 Gb/s) and technology parameters, electrical signaling paths are limited to about 1 m in circuit boards and 10 m in cables. At longer distances, either the signaling rate must be reduced or repeaters inserted to overcome attenuation.

[0028] Historically, the high cost of optical signaling limited its use to very long distances or applications that demanded performance regardless of cost. Recent advances in silicon photonics and their application to active optical cables such as Intel Connects Cables and Luxtera Blazar have provided designers with economical optical interconnects. These active optical cables have electrical connections at either end and electrooptical and optoelectrical modules integrated into the cable itself.

[0029] Although optical cables have a higher fixed cost, their ability to transmit data over long distances at several times the data rate of copper cables results in a lower cost per unit distance than electrical cables. Based on the data available using current technologies, the break-even point is at 10 m. For distances shorter than 10 m, electrical signaling is less expensive. Beyond 10 m, optical signaling is more economical. The Dragonfly topology proposed here exploits this relationship between cost and distance. By reducing the number of global cables, it minimizes the effect of the higher fixed overhead of optical signaling, and by making the global cables longer, it maximizes the advantage of the lower per-unit cost of optical fibers.

[0030] To show an example Dragonfly network topology, the following symbols are used in the description of the dragonfly topology and in example routing algorithms presented later:

[0031] N Number of network terminals

[0032] p Number of terminals connected to each router

[0033] a Number of routers in each group

[0034] k Radix of the routers

[0035] k_{eff} Effective radix of the group (or the virtual router)

[0036] h Number of channels within each router used to connect to other groups

[0037] g Number of groups in the system

[0038] q Queue depth of an output port

[0039] qvc Queue depth of an individual output VC

[0040] H Hop count

[0041] Outi Router output port i

[0042] The Dragonfly topology is a hierarchical network with three levels, as shown in FIG. 1: routers (104, 105, and 106), groups (101, 102, and 103), and system. At the router level, each router has connections to p nodes, a-1 local channels—to other routers in the same group—and h global channels—to routers in other groups. Therefore the radix (or degree) of each router is defined as $k=p+a+h-1$. A group consists of a routers connected via an intra-group interconnection network formed from local channels, as shown at 101 in FIG. 1. Each group has ap connections to terminals and ah connections to global channels, and all of the routers in a group collectively act as a virtual router with radix $k'=a(p+h)$. This very high radix, $k' \gg k$ enables the system level network to be realized with very low global diameter (the maximum number of expensive global channels on the minimum path between any two nodes). Up to $g=ah+1$ groups ($N=ap(ah+1)$ terminals) can be connected with a global diameter of one. In contrast, a system-level network built directly with radix k routers would require a larger global diameter.

[0043] In a maximum-size ($N=ap(ah+1)$) dragonfly, there is exactly one connection between each pair of groups. In smaller dragonflies, there are more global connections out of each group than there are other groups. These extra global connections are distributed over the groups with each pair of groups connected by at least $ah+1$ g_channels.

[0044] The dragonfly parameters a, p, and h can have any values. However, to balance channel load, the network in this example has $a=2p=2h$. Because each packet traverses two local channels along its route (one at each end of the global channel) for one global channel and one terminal channel, this ratio maintains balance. Because global channels are expensive, deviations from this 2:1 ratio are done in some embodiments in a manner that overprovisions local and terminal channels, so that the expensive global channels remain fully utilized. That is, the network is balanced in such examples so that $a \geq 2h$, $2p \geq 2h$.

[0045] The scalability of a balanced dragonfly is shown in FIG. 2. By increasing the effective radix, the dragonfly topology is highly scalable—with radix-64 routers, the topology scales to over 256k nodes with a network diameter of only three hops. Arbitrary networks can be used for the intra-group and inter-group networks in FIG. 1. In the example presented here, we use a 1-D flattened butterfly or a completely-connected topology for both networks. A simple example of the dragonfly is shown in FIG. 3 with $p=h=2$ (two processing nodes per router and two channels within each router coupled to other groups), $a=4$ (four routers in each group) that scales to $N=72$ (72 nodes in the network) with $k=7$ (radix 7) routers. By using virtual routers, the effective radix is increased from $k=7$ to $k'=16$, as group G_0 of FIG. 3 has eight global connections and eight node connections.

[0046] The global radix, k' , can be increased further by using a higher-dimensional topology for the intra-group network. Such a network may also exploit intra-group packaging locality. For example, a 2-D flattened butterfly is shown in FIG. 4 at 401, which has the same k' as the group shown in FIG. 5 but exploits packaging locality by providing more

bandwidth to local routers. A 3-dimension flattened butterfly is used in FIG. 4 at 402 to increase the effective radix from $k'=16$ to $K'=32$ —allowing the topology to scale up to $N=1056$ using the same $k=7$ router as in FIG. 1.

[0047] To increase the terminal bandwidth of a high-radix network such as a dragonfly, channel slicing can be employed. Rather than make the channels wider, which would decrease the router radix, multiple network can be connected in parallel to add capacity. Similarly, the dragonfly topology in some embodiments can also utilize parallel networks to add capacity to the network. In addition, the dragonfly networks described so far assumed uniform bandwidth to all nodes in the network. However, if such uniform bandwidth is not needed, bandwidth tapering can be implemented by removing inter-group channels among some of the groups.

[0048] A variety of minimal and non-minimal routing algorithms can be implemented using the dragonfly topology. Some embodiments of global adaptive routing using local information lead to limited throughput and very high latency at intermediate loads. To overcome these problems, we introduce new mechanisms to global adaptive routing, which provide performance that approaches an ideal implementation of global adaptive routing.

[0049] Minimal routing in a dragonfly from source node s attached to router R_s in group G_s to destination node d attached to router R_d in group G_d traverses a single global channel and is accomplished in three steps:

[0050] Step 1: If $G_s=G_d$ and R_s does not have a connection to G_d , route within G_s from R_s to R_a , a router that has a global channel to G_d .

[0051] Step 2: If $G_s=G_d$, traverse the global channel from R_a to reach router R_b in G_d .

[0052] Step 3: If $R_b=R_d$, route within G_d from R_b to R_d .

[0053] This minimal routing works well for load-balanced traffic, but results in poor performance on adversarial traffic patterns. To load-balance adversarial traffic patterns, Valiant's algorithm can be applied at the system level—routing each packet first to a randomly-selected intermediate group G_i and then to its final destination d. Applying Valiant's algorithm to groups suffices to balance load on both the global and local channels. This randomized non-minimal routing traverses at most two global channels and requires five steps:

[0054] Step 1: If $G_s=G_i$ and R_s does not have a connection to G_i , route within G_s from R_s to R_a , a router that has a global channel to G_i .

[0055] Step 2: If $G_s=G_i$ traverse the global channel from R_a to reach router R_x in G_i .

[0056] Step 3: If $G_i=G_d$ and R_x does not have a connection to G_d , route within G_i from R_x to R_y , a router that has a global channel to G_d .

[0057] Step 4: If $G_i=G_d$, traverse the global channel from R_y to router R_b in G_d .

[0058] Step 5: If $R_b=R_d$, route within G_d from R_b to R_d .

[0059] To prevent routing deadlock, two virtual channels (VCs) are needed for minimal routing and three VCs are required for non-minimal routing, as shown in FIG. 5. These virtual router assignments eliminate all channel dependencies due to routing. For some applications, additional virtual channels may be required to avoid protocol deadlock—e.g., for shared memory systems, separate sets of virtual channels may be required for request and reply messages.

[0060] A variety of routing algorithms for the dragonfly topology have been evaluated, including:

[0061] Minimal (MIN): The minimal path is taken as described previously.

[0062] Valiant (VAL) [32]: Randomized non-minimal routing as described previously.

[0063] Universal Globally-Adaptive Load-balanced [29] (UGALG, UGAL-L) UGAL chooses between MIN and VAL on a packet-by-packet basis to load-balance the network. The choice is made by using queue length and hop count to estimate network delay and choosing the path with minimum delay. We implement two versions of UGAL.

[0064] UGAL-L—uses local queue information at the current router node.

[0065] UGAL-G—uses queue information for all the global channels in Gs—assuming knowledge of queue lengths on other routers. While difficult to implement, this represents an ideal implementation of UGAL since the load-balancing is required of the global channels, not the local channels.

[0066] Cycle accurate simulations are used to evaluate the performance of the different routing algorithms. We simulate a single-cycle, input-queued router switch but provide sufficient speedup in order to generalize the results and ensure that routers do not become the bottleneck of the network. Packets are injected using a Bernoulli process. The simulator is warmed up under load without taking measurements until steady-state is reached. Then a sample of injected packets is labeled during a measurement interval. The simulation is run until all labeled packets exit the system. Unless otherwise noted, the simulation results are shown for dragonfly of size 1K node using $p=h=4$ and $a=8$ parameters. Simulations of other size networks follow the same trend and are not presented due to space constraints. Single flit (flow control unit) packets are used to separate the routing algorithm from flow control issues such as the use of wormhole or virtual cut-through flow control. The input buffers are initially assumed to be 16 flits deep. The impact of different buffer sizes is also evaluated.

[0067] The different routing algorithms are evaluated using both benign and adversarial synthetic traffic patterns, as shown in FIG. 6. Latency v. offered load is shown for the four routing algorithms, using both uniform random traffic at 601 and adversarial traffic at 602. The use of a synthetic traffic pattern allows us to stress the topology and routing algorithm to fully evaluate the network. For benign traffic such as uniform random (UR), MIN is sufficient to provide low latency and high throughput, as shown at 601 of FIG. 6. VAL achieves approximately half of the network capacity because its load-balancing doubles the load on the global channels. Both UGAL-G and UGAL-L approach the throughput of MIN, but with slightly higher latency near saturation. The higher latency is caused by the use of parallel or greedy allocation where the routing decision at each port is made in parallel. The use of sequential allocation will reduce the latency at the expense of a more complex allocator.

[0068] To test the load-balancing ability of a routing algorithm, we use a worst-case (WC) traffic pattern where each node in group G_i sends traffic to a randomly selected node in group G_{i+1} . With minimal routing, this pattern will cause all nodes in each group G_i to send all of their traffic across the single global channel to group G_{i+1} . Non-minimal routing is

required to load balance this traffic pattern by spreading the bulk of the traffic across the other global channels.

[0069] The evaluation for this WC traffic is shown in FIG. 6 at 602. Because MIN forwards all of the traffic from each group across a single channel, its throughput is limited to $1/ah$. VAL achieves slightly under 50% throughput which is the maximum possible throughput with this traffic. UGAL-G achieves similar throughput as VAL but UGAL-L leads to both limited throughput as well as high average packet latency at intermediate load. In the following section, we show how the indirect nature of adaptive routing on the dragonfly leads to performance degradation. We identify the issues with UGAL-L and present mechanisms that can overcome these problems.

[0070] Adaptive routing on the dragonfly is challenging because it is the global channels, the group outputs, that need to be balanced, not the router outputs. This leads to an indirect routing problem. Each router must pick a global channel to use using only local information that depends only indirectly on the state of the global channels. Previous global adaptive routing methods used local queue information, source queues and output queues, to generate accurate estimates of network congestion. In these cases, the local queues were an accurate proxy of global congestion, because they directly indicated congestion on the routes they initiated.

[0071] With the dragonfly topology, however, local queues only sense congestion on a global channel via backpressure over the local channels. If the local channels are overprovisioned, significant numbers of packets must be enqueued on the overloaded minimal route before the source router will sense the congestion. This results in a degradation in throughput and latency as shown earlier in FIG. 6 at 602.

[0072] A throughput issue with UGAL-L arises due to a single local channel handling both minimal and non-minimal traffic. For example, in FIG. 7, a packet in R1 has a minimal path which uses gc7 and a nonminimal path which uses gc6. Both paths share the same local channel from R1 to R2. Because both paths share the same local queue (and hence have the same queue occupancy) and the minimal path is shorter (one global hop vs two), the minimal channel will always be selected, even when it is saturated. This leads to the minimal global channel being overloaded and the non-minimal global channels that share the same router as the minimal channel being under utilized. With UGAL-G, the minimal channel is preferred and the load is uniformly balanced across all other global channels. With UGAL-L, on the other hand, the non-minimal channels on the router that contains the minimal global channel are under utilized—resulting in a degradation of network throughput.

[0073] To overcome this limitation, we modify the UGAL algorithm to separate the queue occupancy into minimal and nonminimal components by using individual VCs (UGAL-LVC).

```

if (qm vcHm ≤ qnm vcHnm )
    route minimally;
else
    route nonminimally;

```

where the subscript m and nm denote the minimal and non-minimal paths. If the VC assignment of FIG. 5 is used, $qm\ vc=q(V\ C1)$ and $qnm\ vc=q(V\ C0)$.

[0074] When compared, UGAL-LVC matches the throughput of UGAL-G on a WC traffic pattern but for UR traffic, the throughput is limited, with approximately 30% reduction in throughput. For the WC traffic, where most of the traffic needs to be sent non-minimally, UGALLVC performs well since the minimal queue is heavily loaded. However, for load-balanced traffic when most traffic should be sent minimally, individual VCs do not provide an accurate representation of the channel congestion—resulting in throughput degradation.

[0075] To overcome this limitation, we further modify the UGAL algorithm to separate the queue occupancy into minimal and non-minimal components only when the minimal and nonminimal paths start with the same output port. Our hybrid modified UGAL routing algorithm (UGAL-LVC H) is:

```

if (qmHm ≤ qnmHnm && Outm_ = Outnm ) ||
(qm vcHm ≤ qnm vcHnm && Outm = Outnm)
    route minimally;
else
    route nonminimally;

```

[0076] Compared to UGAL-LVC, UGAL-LVC H provides the same throughput on WC traffic pattern but matches the throughput of UGAL-G on UR traffic but resulting in nearly 2× higher latency at an offered load of 0.8, near saturation. For WC traffic, UGAL-L VC H also results in higher intermediate latency compared to UGAL-G.

[0077] The high intermediate latency of UGAL-L is due to minimally-routed packets having to fill the channel buffers between the source and the point of congestion before congestion is sensed. Our research shows that non-minimally routed packets have a latency curve comparable to UGAL-G while minimally-routed packets see significantly higher latency. As input buffers are increased, the latency of minimally-routed packets increases and is proportional to the depth of the buffers. A histogram of latency distribution shows two clear distributions—one large distribution with low latency for the non-minimal packets and another distribution with a limited number of packets but with much higher latency for the minimal packets.

[0078] To understand this problem with UGAL-L, in the example dragonfly group shown in FIG. 7, assume a packet in R1 is making its global adaptive routing decision of routing either minimally through gc0 or non-minimally through gc7. The routing decision needs to load balance global channel utilization and ideally, the channel utilization can be obtained from the queues associated with the global channels, q0 and q3. However, q0 and q3 queue informations are only available at R0 and R2 and not readily available at R1—thus, the routing decision can only be made indirectly through the local queue information available at R1.

[0079] In this example, q1 reflects the state of q0 and q2 reflects the state of q3. When either q0 or q3 is full, the flow control provides backpressure to q1 and q2 as shown with the arrows in FIG. 7. As a result, in steady-state measurement, these local queue information can be used to accurately measure the throughput. Since the throughput is defined as the offered load when the latency goes to infinity (or the queue occupancy goes to infinity), this local queue information is sufficient. However, q0 needs to be completely full in order for q1 to reflect the congestion of gc0 and allow R1 to route packets non-minimally. Thus, using local information

requires sacrificing some packets to properly determine the congestion—resulting in packets being sent minimally having much higher latency. As the load increases, although minimally routed packets continue to increase in latency, more packets are sent non-minimally and results in a decrease in average latency until saturation.

[0080] In order for local queues to provide a good estimate of global congestion, the global queues need to be completely full and provide a stiff backpressure towards the local queues. The stiffness of the backpressure is inversely proportional to the depth of the buffer—with deeper buffers, it takes longer for the backpressure to propagate while with shallower buffers, a much stiffer backpressure is provided. As the buffer size decreases, the latency at intermediate load is decreased because of the stiffer backpressure. However, using smaller buffers comes at the cost of reduced network throughput.

[0081] To overcome the high intermediate latency, we propose using credit round-trip latency to sense congestion faster and reduce latency. In credit-based flow control, illustrated in FIG. 8, credit counts are maintained for buffers downstream. As packets are sent downstream, the appropriate credit count is decremented and once the packet leaves downstream router, credits are sent back upstream and the credit count is incremented. The latency for the credits to return is referred to as credit round-trip latency (tcrt) and if there is congestion downstream, the packet cannot be immediately processed and results in an increase in tcrt.

[0082] Referring to FIG. 8, conventional credit flow control is illustrated at 801. As packets are sent downstream (1), the output credit count is decremented (2) and credits are sent back upstream (3). This scheme is modified as shown at 802 to use credit round trip latency to estimate congestion in the network. In addition to the output credit count being decremented (2), the time stamp is pushed into the credit time queue, denoted CTQ. Before sending the credit back upstream (4), the credit is delayed (3), and when downstream credits are received (5), the credit count is updated as well as the credit round trip latency tcrt.

[0083] The value of tcrt can be used to estimate the congestion of global channels. By using this information to delay upstream credits, we stiffen the backpressure and more rapidly propagate congestion information up stream. For each output O, tcrt(O) is measured and the quantity td(O)=tcrt(O)–tcrtO is stored in a register. Then, when a flit is sent to output O, instead of immediately sending a credit back upstream, the credit is delayed by td(O)–min [td(o)]. The credits sent across the global channels are not delayed. This ensures that there is no cyclic loop in this mechanism and allows the global channels to be fully utilized.

[0084] The delay of returning credits provides the appearance of shallower buffers to create a stiff backpressure. However, to ensure that the entire buffer gets utilized and there is no reduced throughput at high load, the credits needs to be delayed by the variance of td across all outputs. We estimate the variance by finding min [td(o)] value and using the difference. By delaying credits, the upstream routers observe congestion at a faster rate (compared to waiting for the queues to fill up) and leads to better global adaptive routing decisions.

[0085] The UGAL-L routing algorithm evaluation using credit latency (UGAL-LCR) is investigated for both WC and UR traffic using buffers of depth 16 and 256. UGAL-LCR leads to significant reduction in latency compared to UGALL and approaches the latency of UGAL-G. For WC traffic, UGAL-LCR reduces latency by up to 35% with 16 buffers

and up to over 20× reduction in intermediate latency with 256 buffers compared to UGAL-L. Unlike UGAL-L, the intermediate latency with UGAL-LCR is independent of buffer size. For UR traffic, UGAL-LCR provides up to 50% latency reduction near saturation compared to UGAL-LVC H. However, both UGAL-LCR and UGALL VC H fall short of the throughput of UGAL-G with UR traffic because their imprecise local information results in some packets being routed non-minimally.

[0086] The implementation of this scheme results in minimal complexity overhead as the following three features are needed at each router:

- [0087]** tracking credits individually to measure tcrt
- [0088]** registers to store td values
- [0089]** a delay mechanism in returning credits

The amount of storage required for td is minimal as only $O(k)$ registers are required. The credits are often returned by piggybacking on data flits and delaying credits to wait for the transmission of the next data flit upstream is required. The proposed mechanism only requires adding additional delay.

[0090] As for tracking individual credits, credits are conventionally tracked as a pool of credits in credit flow control—i.e., a single credit counter is maintained for each output VC and increments when a credit is received. The implementation of UGAL-LCR requires tracking each credit individually. This can be done by pushing a timestamp on the tail of a queue each time a flit is sent, as shown in FIG. 17(b) with the use of a credit timestamp queue (CTQ), and popping the timestamp off the head of the queue when the corresponding credit arrives. Because flits and credits are 1:1 and maintain ordering, the simple queue suffices to measure round-trip credit latency. The depth of the queue needs to be proportional to the depth of the data buffers but the queue size can be reduced to utilize imprecise information to measure congestion—e.g., by having a queue which is only $\frac{1}{4}$ of the data buffer size, only one of four credits are tracked to measure the congestion.

[0091] The cost of a dragonfly topology also compares favorably to a flattened butterfly, as well as to other topologies. The flattened butterfly topology reduces network cost of a butterfly by removing intermediate routers and channels. As a result, the flattened butterfly reduces cost by approximately 50% compared to a folded-Clos on balanced traffic. The dragonfly topology extends the flattened butterfly by increasing the effective radix of the routers to further reduce the cost and increase the scalability of the network.

[0092] A comparison of dragonfly and flattened butterfly networks of 64k nodes shows that a flattened butterfly uses 50% of the router ports for global channels, while a dragonfly uses 25% of the ports for global connections. The flattened butterfly requires two additional dimensions, while the dragonfly is a single dimension. In addition, the dragonfly provides better scalability because the group size can be increased to scale the network whereas scaling the flattened butterfly requires adding additional dimensions. With the hop count nearly identical, the dragonfly trades off longer global cables for smaller number of global cables required to provide a more cost-efficient topology better matched to emerging signaling technologies.

[0093] The dollar cost of a dragonfly also compares favorably to a flattened butterfly for networks larger than 1k nodes, showing approximately a 10% savings for up to 4k nodes, and approximately a 20% cost savings relative to flattened butterfly topologies for more than 4k nodes as the dragonfly has

fewer long, global cables. Folded Clos and 3-d torus networks suffer in comparison, because of the larger number of cables needed to support high network diameters. For a network of only 1k nodes, the dragonfly is 62% the cost of a 3-d torus network and 50% that of a folded Clos network. This reduction in network cost is directly correlated to a reduction in network power consumed, which is a significant advantage for large networks as well as for installations that are desirably environmentally friendly.

[0094] The example embodiments of a dragonfly network presented here show how use of a group of routers as a virtual router can increase the effective radix of a network, and hence reduce network diameter, cost, and latency. Because the dragonfly topology reduces the number global cables in a network, while at the same time increasing their length, the dragonfly topology is particularly well suited for implementations using emerging active optical cables—which have a high fixed cost but a low cost per unit length compared to electrical cables. Using active optical cables for the global channels, a dragonfly network reduces cost by 20% compared to a flattened butterfly and by 52% compared to a folded Clos network of the same bandwidth.

[0095] Various embodiments of dragonfly networks described here also comprise two new variants of global adaptive routing that overcome the challenge of indirect adaptive routing presented by the dragonfly. A dragonfly router will typically make a routing decision based on the state of a global channel attached to a different router in the same group. Conventional global adaptive routing algorithms that use local queue occupancies to infer the state of this remote channel give degraded throughput and latency. We introduce the selective use of virtual channel discrimination to overcome the bandwidth degradation. We also introduce the use of credit round-trip latency to both sense and signal channel congestion. The combination of these two techniques gives a global adaptive routing algorithm that approaches the performance of an ideal algorithm with perfect knowledge of remote channel state.

[0096] Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiments shown. This application is intended to cover any adaptations or variations of the example embodiments of the invention described herein. It is intended that this invention be limited only by the claims, and the full scope of equivalents thereof.

1. A multiprocessor computer system comprising a dragonfly processor interconnect network, the dragonfly processor interconnect network comprising:
 - a plurality of processor nodes;
 - a plurality of routers, each router directly coupled to a plurality of terminal nodes, the routers coupled to one another and arranged into a group,
 - and a plurality of groups of routers, such that each group is connected to each other group via at least one direct connection.
2. The multiprocessor computer system of claim 1, wherein each group acts as virtual router with radix of approximately 2 times the square root of the number of nodes in the network.
3. The multiprocessor computer system of claim 1, wherein the virtual radix of each group is the product of the

number of routers in each group times the sum of the number of processor nodes connected to each router plus the number of global channels

4. The multiprocessor computer system of claim 1, wherein the number of routers per group is equal to twice the number of processor nodes per router, and wherein the number of processor nodes per router is equal to the number of channels per router connected to other groups.

5. The multiprocessor computer system of claim 1, wherein the number of routers in a group is greater than twice the number of global channels per router.

6. The multiprocessor computer system of claim 1, wherein the number of processor nodes per router is greater than the number of global channels per router.

7. The multiprocessor computer system of claim 1, wherein the routers within a group are connected via a flattened butterfly network.

8. The multiprocessor computer system of claim 1, wherein the routers route data using selective virtual channel discrimination.

9. The multiprocessor computer system of claim 1, wherein the routers route data using credit round-trip latency as an indicator of channel congestion.

10. A method of operating a multiprocessor computer system, comprising:

communicating a message from a processor node to a router, the router coupled to a plurality of processor nodes;

communicating the message between two or more routers, the routers coupled to one another and arranged into a group, and

communicating the data between two groups of routers, such that each group is connected to each other group via at least one direct connection.

11. The method of operating a multiprocessor computer system of claim 1, wherein each group acts as virtual router

with radix of approximately 2 times the square root of the number of nodes in the network.

12. The method of operating a multiprocessor computer system of claim 1, wherein the virtual radix of each group is the product of the number of routers in each group times the sum of the number of processor nodes connected to each router plus the number of global channels

13. The method of operating a multiprocessor computer system of claim 1, wherein the number of routers per group is equal to twice the number of processor nodes per router, and wherein the number of processor nodes per router is equal to the number of channels per router connected to other groups.

14. The method of operating a multiprocessor computer system of claim 1, wherein the number of routers in a group is greater than twice the number of global channels per router.

15. The method of operating a multiprocessor computer system of claim 1, wherein the number of processor nodes per router is greater than the number of global channels per router.

16. The method of operating a multiprocessor computer system of claim 1, wherein the routers within a group are connected via a flattened butterfly network.

17. The method of operating a multiprocessor computer system of claim 1, wherein the routers route data using selective virtual channel discrimination.

18. The method of operating a multiprocessor computer system of claim 1, wherein the routers route data using credit round-trip latency as an indicator of channel congestion.

19. A multiprocessor computer system, comprising a Dragonfly processor interconnect network.

20. A method of communicating data between processing nodes in a multiprocessor computer system, comprising routing the data over a Dragonfly processor interconnect network.

* * * * *