



(51) International Patent Classification:

H04L 12/911 (2013.01) H04L 12/927 (2013.01)

H04L 12/931 (2013.01) H04L 12/741 (2013.01)

H04L 12/801 (2013.01) H04L 29/12 (2006.01)

(21) International Application Number:

PCT/US2020/024260

(22) International Filing Date:

23 March 2020 (23.03.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/852,273 23 May 2019 (23.05.2019) US

62/852,203 23 May 2019 (23.05.2019) US

62/852,289 23 May 2019 (23.05.2019) US

(71) Applicant: **CRAY INC.**, [US/US]; 901 5th Ave., STE 1000, Seattle, Washington 98164-2008 (US).

(72) Inventors: **BEECROFT, Johnathan P.**; 901 5th Ave., STE 1000, Seattle, Washington 98164 (US). **TURNER, Edward J.**; 901 5th Ave., STE 1000, Seattle, Washington 98164 (US).

(74) Agent: **FEBBO, Michael** et al.; HEWLETT PACKARD ENTERPRISE, 3404 E. Harmony Road Mail Stop 79, Fort Collins, Colorado 80528 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) Title: SYSTEM AND METHOD FOR DYNAMIC ALLOCATION OF REDUCTION ENGINES




FIG. 1

(57) Abstract: A switch equipped with a reduction engine capable of being dynamically allocated in a network is provided. During operation, the reduction engine can be dynamically armed based on a multicast frame. As a result, the network can facilitate an efficient and scalable environment for high performance computing.

**Declarations under Rule 4.17:**

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

- *with international search report (Art. 21(3))*

# SYSTEM AND METHOD FOR DYNAMIC ALLOCATION OF REDUCTION ENGINES

5                   **Inventors:** Jonathan P. Beecroft and Edward J. Turner

## BACKGROUND

### Field

10           **[0001]** This is generally related to the technical field of networking. More specifically, this disclosure is related to systems and methods for facilitating dynamic allocation of reduction engines in a network.

### Related Art

15           **[0002]** As network-enabled devices and applications become progressively more ubiquitous, various types of traffic as well as the ever-increasing network load continue to demand more performance from the underlying network architecture. For example, applications such as high-performance computing (HPC), media streaming, and Internet of Things (IOT) can generate different types of traffic with distinctive characteristics. As a result, in addition to conventional network performance metrics such as bandwidth and delay, network architects  
20           continue to face challenges such as scalability, versatility, and efficiency.

## SUMMARY

**[0003]** A switch equipped with a reduction engine capable of being dynamically allocated in a network is provided. During operation, the reduction engine can be dynamically armed  
25           based on a multicast frame. As a result, the network can facilitate an efficient and scalable environment for high performance computing.

## BRIEF DESCRIPTION OF THE FIGURES

**[0004]** FIG. 1 shows an exemplary network.  
30           **[0005]** FIG. 2 shows an exemplary multicast tree for a reduction process.  
             **[0006]** FIG. 3A shows a flow chart of an exemplary reduction process.  
             **[0007]** FIG. 3B shows a flow chart of an exemplary reduction operation by a reduction engine.  
             **[0008]** FIG. 4 shows an example where one leaf endpoint is late joining the reduction  
35           process.

[0009] FIG. 5A shows an example where one leaf endpoint fails to supply a contribution because of an error.

[0010] FIG. 5B shows a flow chart of an exemplary timer-based reduction process.

[0011] FIG. 6 shows an example where a reduction engine on a leaf switch is unavailable.

5 [0012] FIG. 7 shows exemplary reduction operations.

[0013] FIG. 8 shows a set of MINMAXLOC operands that can be used in a reduction process.

[0014] FIG. 9 shows rounding modes that can be used in a reduction process.

[0015] FIG. 10 shows a Portals-formatted reduction frame.

10 [0016] FIG. 11 shows a reduction header.

[0017] FIG. 12 shows the endianness of operands that can be used for MINMAXLOC reproducible sum operators in a reduction process.

[0018] FIG. 13 shows exemplary reduction result codes

15 [0019] FIG. 14 shows an example where a Portals packet can be prepended with an Ethernet header.

[0020] FIG. 15 shows an exemplary switching system that facilitates a reduction engine.

[0021] In the figures, like reference numerals refer to the same figure elements.

## DETAILED DESCRIPTION

20 [0022] Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present disclosure. Thus, the present invention is not limited to the embodiments shown.

[0023] Embodiments of the present invention solve the problem of accommodating a  
25 large number of computing endpoints in a network by providing a reduction engine that can be dynamically configured, which allows traffic resulting from large-scale computing to be reduced in a timely, flexible, and scalable manner. Allocation and management of any shared resource within the body of a network can be difficult, especially if errors occur while the unit is processing data. The systems and methods described herein can significantly simplify the  
30 allocation, deallocation, and error handling of a dynamically allocated switch/router resource.

[0024] In general, a network can support thousands of users. If a function is provided, and expected to be used by any user, it is important to manage the resource efficiently. One approach can be to use a system call to a network-based server that has been authorized to manage the function. Although this may appear to be a simple solution, in practice the  
35 management could become complicated, especially if the function provided is widely distributed

throughout the network and thousands of users may be trying to gain access. The time it takes to setup an operation for a single use can run into many seconds, and a similar amount of time may be required to release the function after use. Any error condition may also require significant support in both software development and real time analysis during the error condition.

5           **[0025]** The type of function being used may only be active for a few microseconds or even a few nanoseconds. Even if the function is repeatedly used by the same application, the setup and teardown cost could dwarf the amount of time the function is being useful. For computing features with such an enormous overhead, using software to gain and release access to such functions may not be able to accommodate a large number of users.

10           **[0026]** Embodiments of the present invention can provide reasonably fair access to all users of the network while reducing the setup cost tiny and ensure the resource can be released quickly on a successful completion of the operation. It can also ensure the resource is released reasonably quickly (e.g., a few milliseconds) when an error occurs, without the need for any management software intervention.

15           **[0027]** Specifically, a reduction engine can be provided within a switch. The reduction engine can take packets from a number of endpoints and combine them to generate a single packet that can be returned to a node. The reduction engine can also perform a synchronization function, often referred to as a barrier, or can perform some mathematical function that combines or sorts the values provided by the endpoints into a single value. By placing the reduction engine  
20 within the body of the network, the latency, i.e., the time it takes to complete the operation, can be reduced by an order of magnitude because typically a single round of communication across the network is usually sufficient to complete the whole reduction.

**[0028]** The reduction process can use a multicast session, issued from a reduction root node's edge port and sent to all the edge leaf ports, to setup or arm each of the reduction engines  
25 within the network. Each port of the switches within the network can have an instance of the reduction engine. The arming multicast (i.e., the multicast setup packet) can start at the root node's ingress edge port. When the setup packet is received, it can arm the local reduction engine associated with the corresponding port. The packet can then be multicast to a number of output ports where it is forwarded to the output's link partner ingress port, which can reside on  
30 another switch. The downstream switch can further multicast the setup packet to a set of output ports while at the same time arming an instance of a reduction engine associated with the ingress port.

**[0029]** The above process can repeat, arming reduction engines along the multicast data path until the setup packet arrives at the egress edge ports of the leaf switches where it is passed  
35 to the compute node. At this point, all the reduction engines can be armed and ready to receive

the reduction packets, which can travel upstream along the multicast tree and be reduced to a single packet that represents a reduced result.

[0030] After a computation operation, the leaf nodes can be ready to inject their result packets back into the network. They can do so in a way that causes the packet to retrace on the reverse path taken by the original multicast packet. This ensures the result packets can each be intercepted by the now armed reduction engines where the reduction function can take place.

[0031] The multicast tree can be traversed by the result packets in the reverse direction through the network. This reduction process does not require any software intervention, other than setting up the initial multicast tree. Modern switching devices typically can accommodate a large number of separate multicast trees, which allows many reduction configurations to be simultaneously configured in a network. As a result, the setup and teardown cost can be significantly amortized and parallelized, which allows the reduction function to be scaled to a large number of users.

[0032] It is also possible to add a count value to each reduction packet to represent the number of inputs used to construct the reduction result held within the packet. This allows the acceleration provided by a reduction engine to be skipped, if necessary, without affecting the reduction function. This is possible because the receiving node is then given enough information to complete the operation itself.

[0033] A timeout mechanism can also be added to the reduction engines to ensure the reduction engine resource can eventually become free, even in the presence of errors. If an error occurs, or if the reduction is not able to complete because one of the inputs to the reduction function is not present or is delayed for some reason, the timeout can ensure that the resource is released with the available input information that can be used for the reduction computation, up to that point. The root node of the reduction can receive this partial result and recognize that this result is not complete. The root node can optionally wait for the missing result to arrive without blocking the shared reduction resource.

[0034] FIG. 1 shows an exemplary network. In this example, a network 100 of switches, which can also be referred to as a “switch fabric,” can include switches 102, 104, 106, 108, and 110. Each switch can have a unique address or ID within switch fabric 100. Various types of devices and networks can be coupled to a switch fabric. For example, a storage array 112 can be coupled to switch fabric 100 via switch 110; an InfiniBand (IB) based HPC network 114 can be coupled to switch fabric 100 via switch 108; a number of end hosts, such as host 116, can be coupled to switch fabric 100 via switch 104; and an IP/Ethernet network 118 can be coupled to switch fabric 100 via switch 102. In general, a switch can have edge ports and fabric ports. An edge port can couple to a device that is external to the fabric. A fabric port can couple to another

switch within the fabric via a fabric link. Typically, traffic can be injected into switch fabric 100 via an ingress port of an edge switch, and leave switch fabric 100 via an egress port of another (or the same) edge switch. An ingress link can couple a network interface controller (NIC) of an edge device (for example, an HPC end host) to an ingress edge port of an edge switch. Switch fabric 100 can then transport the traffic to an egress edge switch, which in turn can deliver the traffic to a destination edge device via another NIC.

**[0035]** In one embodiment, each port of a switch can include a reduction engine that is used to accelerate reduction operations. Reductions can be performed using a multicast tree. Each reduction engine in the multicast tree can be armed by a reduction arm frame sent by a root switch through the multicast tree. After receiving the reduction arm frame, leaf nodes of the multicast tree can send reduction data frames containing their contributions up the multicast tree to the root node. Each reduction engine in the tree can intercept the reduction data frames and perform reduction on them. When a reduction engine receives the expected number of contributions or times out, it can forward the reduced result up the multicast tree. The root node may receive a single, fully reduced data frame, or, if any reduction engine times out, it may receive multiple, partially reduced data frames. In either case, the root node can complete the reduction, incorporating its own contribution. The final result of the reduction can then be sent down the multicast tree to leaf nodes. The result frame can carry another round of reduction arming instruction, which can then re-arm the reduction engines at the same time.

**[0036]** The reduction engine can reduce latency in critical network operations including reduce, all-reduce, and barrier. Reduction operations can be performed over a spanning tree embedded within the network. FIG. 2 shows an exemplary multicast tree for a reduction process. In this example, a multicast tree for the reduction process can include a root endpoint 202, a root switch 204, a number of leaf switches such as leaf switch 206, and a number of leaf endpoints such as endpoint 208. Root switch 204 is responsible for initiating the multicast tree for the reduction process. Each switch can include a reduction engine that can be armed when the multicast session is setup. The leaf endpoints can inject frames, which can be combined as they flow up the tree, with the result being delivered to a process running at the root of the tree. As described below, the root process may need to complete the reduction in software. This is the ready phase of a reduction. The result of a reduction can be then multicast back down the tree to processes at the leaf endpoints and the reduction engines can be re-armed, ready for the next round reduction. This is the multicast phase of a reduction process.

**[0037]** The multicast phase of a reduction process can provide synchronization for a barrier operation, during which no data is required and a null reduction operation is used. Each node can join the reduction tree and wait for the result. When the root node receives the result, it

can then issue a multicast down the reduction tree. In one embodiment, no endpoint is allowed to leave the barrier before all endpoints have entered.

[0038] Reduction engines can be provided on the output side of each link. They can operate on data held in the reduction buffers. In one embodiment, each reduction engine can support eight active reduction trees. Other numbers of reduction trees can also be supported. The reduction engines can perform on-the-fly combining of data frames. The reduction engines are armed during the multicast phase. They can combine upstream frames for a given amount of time. The reduction engine can be disarmed either when the current operation has been completed, or after a timeout period. In the event of a reduction timing out, any partial results can be forwarded up the tree towards the root. The purpose of the timeout is to ensure that no reduction state remains in the event of error, device failure, or frame loss within the reduction tree.

[0039] FIG. 3A shows a flow chart of an exemplary reduction process. During operation, a root process first initializes the reduction tree (operation 302). In HPC programming models, initialization can be a collective operation involving a number of processes that are to participate in a reduction. One process, which in this case can be the root process, can communicate with the network management software to create a spanning tree (the multicast tree), which can be represented by a multicast address. The network can use a multicast protocol to establish the multicast tree topology, and store the forwarding information in a data structure, such as a multicast table. This data structure typically stores topological and forwarding information, such as for a given multicast address, what output ports should a multicast packet be forwarded to. The root process can then arm the reduction engines in the spanning tree by sending a frame to the multicast address (operation 304). Other processes can wait until they receive this frame. Once this frame has reached all the participating processes, the reduction tree is now ready for use.

[0040] Subsequently, the participating processes can perform the computation task that results in their contribution to the reduction operation. Processes other than the root process can each construct a reduction frame and send it to the multicast address of the reduction tree (operation 306). The reduction engines residing in the switches participating in the reduction tree can perform reduction on the received frames and each send a reduced frame upstream toward the root switch of the reduction tree (operation 308). The root process can consume the data reduction frames. It can receive the contributions from the leaf nodes in one or more data reduction frames and complete the ready phase by performing the reduction operation to these frames, including its own contribution (operation 310). Optionally, the root process can then determine whether the computation task is complete (operation 312). If it complete, the root

process can send the result to the multicast address and release the reduction engines (operation 316). If the computation task is not complete, the root process subsequently constructs a reduction frame containing the result and sends it to the multicast address, which in turn can re-arm all the reduction engines in the reduction tree (operation 314). This operation can prepare the reduction engines for the next round of reduction. A similar reduction process can then be repeated until the computation task is complete.

[0041] The root node, or more generally a process on the root node, can perform a special role. It first completes the reduction process. As described later, loops are usually not allowed in the multicast tree; hence the root typically does not send its own contribution to itself. Assuming that the reduction engine at the root node of the reduction tree is able to accumulate all of the contributions from the leaf nodes before timing out, the root node can receive a single data reduction frame from the leaf nodes. In this case, the root node can combine this result with its own contribution. On the other hand, if the reduction engine at the root node times out or cannot be allocated, the root node may receive a number of data reduction frames that are to be combined in software. Once the root node has computed the final reduction, it can multicasts this result to the leaf nodes. The root process can also have additional responsibilities in terms of handling errors.

[0042] FIG. 3B shows a flow chart of an exemplary reduction operation by a reduction engine. During operation, a reduction engine residing on a switch can first receive a barrier frame sent to the multicast address (operation 322). Note that the barrier frame is the initial frame that is used to arm the reduction engines for a reduction tree for the first time. After receiving the barrier frame, the reduction engine can record the parent node (i.e., the switch from which the barrier frame is received), and perform a lookup in the multicast table to determine the downstream node to which the barrier is to be forwarded (operation 324). In addition, the reduction engine also sets a wait count, which corresponds to the number of contributions from endpoints that are expected to be received based on the multicast table entry.

[0043] Subsequently, the reduction engine forwards the barrier frame to the child nodes (operation 326). As the barrier frame travels down the reduction tree, all the reduction engines participating in this reduction tree are armed. As a result, the reduction engine at the local switch begins to receive reduction frames returned from the child nodes or endpoints (operation 328). Next, the reduction engine can aggregate the contributions and forward a reduction frame to the parent node (operation 330). At this point, the reduction engine is now ready for reduction operation. Next, the reduction engine can receive a result frame from the root node, which arms all the reduction engines for a reduction operation which involves actual data.

[0044] The examples shown in this description assume that one process per node contributes to the reduction, which is not required. There can be multiple contributions per node. In some embodiments, a local shared memory reduction and a network reduction can both be performed. Furthermore, the reduction engine described herein can support multiple concurrent non-blocking reduction operations on the same reduction tree.

[0045] The computational operations supported by a reduction engine can include, but are not limited to:

- Null (i.e., the barrier operation which does not involve any payload data);
- MIN, MAX, and SUM operations on integer or floating point data types;
- MINMAXLOC operation (which returns the locations of minimum and maximum values found in an array) on integer or floating point values and integer indices;
- Bitwise AND, OR, and XOR operation on integer data types;
- Reproducible sum operations on floating point data types.

[0046] The data types supported by a reduction engine can include, but are not limited to 64-bit integer and 64-bit IEEE 754 floating point.

[0047] In one embodiment, the MINMAXLOC operator can follow Message Passing Interface (MPI) conventions for MINLOC and MAXLOC operators when the values being compared are equal. In one embodiment, the lower of the two index values is returned.

[0048] For compatibility with a commonly used modern instruction set, rounding modes and exception behavior can follow the definitions in the Advanced RISC Machine (ARM) Architecture Reference Manual, ARMvS. For example, if any operand of a floating point operation is not a number (NaN), the result can be a quiet NaN with its sign=0.

[0049] The reproducible sum and MINMAXLOC operators can use one operand per endpoint. Other reductions can be performed on four 64-bit operands at a time with the same operation being applied to each of the operands.

[0050] The sum of a set of IEEE floating point values may depend on the order in which the operands are added. This can be an important issue when a reduction includes operands of widely varying magnitudes. The publication “Efficient Reproducible Floating Point Reduction Operations on Large Scale Systems,” available at [https://bebop.cs.berkeley.edu/reproblas/docs/talks/SIAM\\_AN13.pdf](https://bebop.cs.berkeley.edu/reproblas/docs/talks/SIAM_AN13.pdf) describes one technique that can be used to achieve the desired level of precision for a given number of elements.

[0051] A deterministic reduction can be performed using a global maximum followed by a global sum using standard floating point arithmetic. A single global sum using integer arithmetic can also be used. With the second approach to reduction, the host software is to perform the same operation when multiple contributions are delivered to the root node.

[0052] In general, each reduction engine can support multiple, independent reduction trees, each identified by a globally unique multicast address. Each point in the tree can be initialized with a local wait count value denoted as *rt\_waitcount*. This count value is normally equal to the number of endpoints beneath that stage of the tree (i.e., the number of children nodes of a given node in the tree).

[0053] Reduction trees can be initialized by creating an entry in a multicast table which specifies the wait count and the set of output ports. This static state, which varies between locations in the tree, can be initialized by the management software in the same way as a multicast tree.

[0054] A single multicast address can be used for each reduction tree. At a parent port, the multicast table entry can specify the set of child ports. At each of the child ports, the multicast table entry can specify the parent port, i.e., the reverse path pointing back towards the root node. In general, loops are not allowed within a reduction tree. Unlike typical multicast entries, where any member of the multicast group is able to multicast to all other members of the multicast group, the multicast entries set up for reduction are one-sided and only the reduction root is able to multicast to all members of the multicast set. When any other member of the reduction tree sends a frame to the multicast address, this frame is only forwarded back to root node of the reduction tree. In addition, the forwarding of this frame typically follows exactly the reverse of the downstream multicast path from the root node. This forwarding mechanism guarantees reduction frames can be correctly intercepted by the reduction engines that have been set up for them.

[0055] In one embodiment, one or more fields in a frame's header can be used as a protection key to ensure that all contributors to a given reduction are from the same application or service. For example, a virtual network identifier (VNI) field from the frame header can be used as a protection key. In addition, a frame's reduction header can contain a 32-bit cookie. All frames in a reduction can be required to have the same protection key and cookie as the frame used to arm all the reduction engines in the same reduction tree.

[0056] As mentioned above, reduction trees are armed before they can be used. A multicast session can be used to arm the tree. In a global reduction process, the multicast phase that distributes the result of a reduction can re-arm the tree. In one embodiment, a *reduction\_arm* request can include the state that is constant for all points in the reduction tree. Reduction operations on a given tree can be identified by their multicast address, a cookie *rt\_cookie*, and a sequence number *rt\_seqno*. All contributors can be required to provide the same protection key (which can be the VNI value), cookie value, and the correct sequence number. The reduction engine can confirm that these conditions are met. The cookie value can be used to help prevent

accidental or malicious interference in the reduction. It may be a random value generated by the root process.

[0057] The process of arming a reduction tree can create a dynamic state in the reduction engines for a given tree. The wait count value can be copied from the multicast table, which indicates the number of output ports for a given multicast address. A timeout value can be determined by comparing the value of the *rt\_waitcount* value with values programmed by the management software. The protection key, cookie value, and the sequence number can be copied from the multicast frame. A local counter *rt\_count*, which tracks the number of received reduction frames, can be initialized to zero.

[0058] All contributions to a given reduction can specify the same reduction operation, which can be identified by an *rt\_op* value. The reduction hardware can generate an error if frames with the same sequence number specify different operations.

[0059] Partial result frames can include a count of the accumulated number of contributions. The leaf endpoints can inject frames with a count of one. The reduction engine can increment the local counter by the partial counts from each frame as it performs the reduction operation. The reduction operation is complete in a given reduction engine when the local count reaches the wait count. On completion of a reduction or expiry of the timeout, the reduction engine forwards the partial result and frees up the dynamic state for the reduction tree. The static state remains in the multicast table until the multicast table entry is deleted. The reduction tree must be rearmed before it can be used again.

[0060] The result of the reduction is completed by a process at the root node. In a global reduction the result can be distributed to the leaf nodes using a multicast down the reduction tree. This operation can also re-arm the tree. In one embodiment, the system can supply both the sequence number for the result being distributed, *rt\_resno*, and the sequence number for the reduction being armed, *rt\_seqno*. A management software can increment the sequence number from one reduction to the next. In normal operation, *rt\_seqno* is typically one higher than *rt\_resno* modulo the size of the counter, which may not be the case in the event of error. For upstream reduction data frames, *rt\_resno* does not need to be set by the management software and hence can be ignored by the hardware. The reduction engine can set *rt\_resno* equal to *rt\_seqno* when sending frames upstream.

[0061] In the example shown in FIG. 2, the reduction tree has a branching ratio of four. Endpoints such as leaf endpoint 208 supply their contributions with a count of one. Each first stage switch in this example, such as switch 206, has a wait count of four, which corresponds to the four leaf endpoints connected to each leaf switch. These switches can combine frames from their four children. Partial result frames with a contribution count of four are forwarded up the

tree to the second stage switch, which in this example is switch 204. Switch 204 can have a wait count of 16. This second stage switch then applies the reduction operator to four frames (one from each of its children) and forwards a result frame with count 16 to root endpoint 202. In practice, the multicast tree may not be completely balanced, where some leaf nodes can have different contribution counts than others, and some later-stage switches can have different contribution counts from other switches at the same level.

[0062] Note that the reduction engine in each switch can accelerate a component of the reduction, which may not always be necessary for proper functionality. A reduction arm command may be unable to allocate a descriptor because perhaps all of the descriptors are busy, or a reduction descriptor may have timed out before all of the results are received. In either case, data frames for this reduction may fail to find a matching descriptor and then be forwarded along the multicast path. A reduction engine in a switch higher in the multicast tree may reduce these frames or they may reach the root where they can be reduced in software.

[0063] In general, a single barrier or reduction operation proceeds with each node other than the root node providing a contribution and then waiting for a result to be returned from the root node. The root node gathers contributions, completes the reduction, and multicasts the result. The sequence number is incremented from one such reduction to the next. It may be desirable to pipeline multiple reductions over the same set of nodes at the same time, for example, to offload progression of non-blocking reductions or to increase bandwidth on multi-element floating point reductions. Software can perform multiple concurrent reductions on the same tree by distinguishing such reductions using high bits of the multicast address. For example, bits 15-3 of the multicast format Destination Fabric Address (DFA), which in one embodiment is used as an inter-switch address for routing traffic within a switch fabric, can distinguish 8K multicast trees. Bits 20-18 of the same multicast address can then be used to distinguish up to 8 reductions being performed concurrently on the same tree. When a reduction engine forms part of more than one reduction tree, probability of contention can increase when multiple reductions are performed concurrently. As a result, more of the reduction operations may be performed higher in the tree.

[0064] In one embodiment, a reduction engine can use a timer to limit the amount of time it spends waiting for contributions. As a result, reduction operations may time out. On expiry of the timer, the partial result of a reduction can be sent up the tree towards the root endpoint. Any further data frames associated with a reduction that has timed out can also be sent up the tree.

[0065] FIG. 4 shows an example where one leaf endpoint 402 (shown in gray) is late joining the reduction process. All the rest leaf endpoints send packets with a count of one. One of the first stage switches, switch 404 (shown in gray) has received frames from three of its

children when the timeout expires. It can then forward a result frame with a count of three up the tree. Sometime later, endpoint 402 supplies its contribution to the reduction. This frame is forwarded up the tree. In this example, second stage switch 406 accumulates five frames, three with counts of four, one with a count of three (from switch 404), and the late frame with a count of one from endpoint 402. Switch 406 subsequently sends its result to root endpoint 408 with a count of 16. In one embodiment, second stage switch 406 can have a longer timeout period to accommodate the late arrival. If all switches had the same timeout values, second stage switch 406 may forward two frames to root endpoint 408, one with a count of 15 and the late frame from endpoint 402 with a count of one. The root endpoint can then complete the reduction.

[0066] In one embodiment, the value for a reduction timeout can be set based on the expected time between reduction operations plus twice the expected variation in arrival time. High timeout values (seconds) do not alter the error free operation, but may delay the arrival of the partial results in the event of error. High timeout values can also cause reduction engine resources to be tied up for longer in the case of a dropped frame or the delayed arrival of a partial result. On the other hand, low timeout values may cause problems with scalability.

[0067] The purpose of the reduction engine is to accelerate latency-sensitive operations, e.g., those in which all processes arrive at a reduction or barrier at approximately the same time. Where there is significant load imbalance and one or more endpoints arrive late, the root node can receive multiple frames. The root node can complete the reduction quickly as the last frame arrives. If the timeout is set correctly (or conservatively), the time spent processing these frames can be small in comparison with the time spent waiting.

[0068] In some systems, such as Exascale systems, errors such as frame drops are usually expected to occur. The reduction mechanism can be designed to still function well in the presence of errors. Two classes of error can be of importance for reduction operations: link errors that cause frames to be corrupted and device errors that cause frames to be dropped. Where available, link level retry can protect against common link errors; therefore, the dominant error case can be expected to be dropped frames arising from a switch or cable failure. Where forward error correction (FEC) is used without link level retry, a small proportion of link errors can result in frame drops. Most such frame drops can occur on bulk data transfer rather than reductions. However, the likelihood of an error causing reduction frames to be lost can increase with the job size.

[0069] The time required to detect errors in reduction operations arising from frame loss can generally be counted in seconds. Under normal operation, this period can be set to be longer than the expected spread of arrival times or the time that nodes might spend in computation while waiting for a non-blocking reduction to complete. After this period, the root process can

reasonably be assumed to be blocked waiting for the reduction to complete. However, there are many examples where processes wait for long periods of time in reduction or barrier operations while other processes complete sequential work. A long time spent in a reduction does not always imply that an error has occurred.

5           **[0070]** FIG. 5A shows an example where one leaf endpoint fails to supply a contribution because of an error. In this example, leaf endpoint 502 experiences an error and does not supply its contribution to the reduction process. As a result, intermediate switch 504 times out and forwards a partial result with a count of three. Root switch 506 again times out (because a total count of 16 is not received before switch 506's timer expires) and forwards a frame with a count  
10 of 15, or perhaps a frame with a count of 12 and a second, subsequent frame with a count of three. In either case, root endpoint 508 then consumes the frame(s) and determines that there has been an error.

**[0071]** FIG. 5B shows a flow chart of an exemplary timer-based reduction process. In this example, a reduction engine first receives a result frame from the root node, which arms the  
15 reduction engine (operation 522). The reduction engine then becomes armed for the reduction tree identified in the result frame, and forwards the result frame to its child nodes in the reduction tree (operation 524). Next, the reduction engine determines whether a sufficient number of reduction contributions have been received (operation 526). This determination is based on the corresponding wait count, which has been set up during the ready phase when the reduction tree  
20 is initiated for the first time. If the wait count has been met, the reduction engine then performs the reduction operation on the received contributions, generates its own reduction frame, and forwards the reduction frame toward the root (operation 532). If the wait count is not satisfied, the reduction engine further determines whether the timer has expired (operation 528). If the timer has not expired, the reduction engine continues to wait for reduction frames to arrive from  
25 child nodes (operation 530). If the timer has expired, the reduction engine then performs reduction operation on the received reduction frames, if any, and forwards its own reduction frame toward the root node (operation 532). If there are any late contribution frames that arrive after the timer expires, the reduction engine can also forward them toward the root node (operation 534).

30           **[0072]** In one embodiment, reliable reductions can be implemented at the transport layer. The network hardware can be designed to accelerate the common case and to free all hardware resources in the case of error. Software can protect against device failure in the reduction tree by performing the same reduction operation over two independent trees. The probability of uncorrelated double error is small. When the root endpoint receives a result (or sufficient frames

to construct the result) from one tree the host software can multicast the result on both trees. The sequence number from the successful operation can be used as the result number.

[0073] In the case where the second tree is left with a partial result, potentially several steps back, or stranded on a congested link or a busy queue, delivery of the reduction multicast frame with the *rt\_arm* bit set can advance the sequence number and clear the state. Frames with an old sequence number can be dropped. One or more frames may still be in flight up the tree.

[0074] Reductions are typically latency sensitive, and not bandwidth intensive. The additional network load created by performing two simultaneous reductions is usually negligible. An added benefit of performing reductions twice, on distinct trees, is that the first result can be used, potentially reducing the time to complete the operation if the network is congested. The drawback to this approach, however, is that twice as many Reduction resources (e.g., multicast table entries and reduction engines/IDs) are used, potentially reducing the number of simultaneous reductions by a factor of two.

[0075] A correlated double error may arise as a result of chassis power failure, but such an error may likely cause node failures as well. In a multi-slice network, software can create reductions trees on different slices. Where there is only a single network slice, trees on the same slice can share a common-chassis switch. The reduction is vulnerable to loss of this switch, but such an error can disconnect the nodes as well.

[0076] Allowing reductions to time out ensures that no reduction state is left behind in the event of error. There does not need to be requirement to issue management requests to search and release hardware resources after a fault. In some embodiments, reduction state can be left in the network in the event of error and fault recovery can be complex.

[0077] In some embodiments, a flexible reduction protocol can be provided where a missing reduction engine, or one where all resources have been consumed, does not render the protocol dysfunctional. A missing early leaf reduction computation can be completed by a reduction engine higher in the tree, as shown in FIG. 6. In this example, the reduction engine on a leaf switch 604 is unavailable for the reduction process. As a result, leaf switch 604 forwards all four contribution frames generated by the four children leaf endpoints to a root switch 606. Meanwhile, other three leaf switches 601, 602, and 603 each send their respective reduction frame (with a count of four) to root switch 606. Root switch can process all these seven frames and forwards a single reduction frame with a count of 16 to root endpoint 608. In general, a reduction computation high in the tree can be completed by a process running on the root node. The acceleration offered by the reduction engines in those cases may be different but the calculated result remains the same.

[0078] If a reduction tree is armed, but no contributions are provided before all of the active reduction engines time out, then all contributions can be forwarded to the root endpoint. The root process can subsequently perform the entire computation. The time required to perform this computation can be much shorter than the timeout period. This reduction scenario does not require acceleration.

[0079] The maximum number of reduction trees that can be supported on a given system can be determined by the product of the number of endpoints that can act as root and the number of active trees supported by a given reduction engine. There is usually limited value in accelerating reductions over small numbers of nodes, because each node can simply send its contribution to the root node. Suppose the goal is to accelerate reductions of sixteen nodes and above. A system might run a thousand or more such jobs, while large systems tend to run a mix of job sizes, which in turn may reduce the number of active reduction trees. A large application might use multiple reduction trees at the same time. In the canonical example, the processes can be arranged in a 2D mesh and reductions are performed over the rows, the columns, and over the entire mesh. For  $P$  processes, the number of active reductions can be twice the square root of  $P$  plus one. For implementation considerations, the multicast table is relatively expensive in terms of die area. In one embodiment, a switch chip can support 8192 multicast addresses, while other numbers are also possible. Note that non-intersecting trees may use the same multicast address.

[0080] In some embodiments, software can decide which reductions to offload. A request can be sent to the network management system to create a reduction tree. If this request fails the network application interface (API) can perform the operation in software. However, running different instances of the same job with software-based reductions or accelerated reductions can potentially lead to performance variation, which is undesirable. The reduction offload strategy can be configured such that running out of reduction trees is unlikely.

[0081] In one embodiment, a 6-bit command field is used to indicate different reduction operations, as shown in FIG. 7. This 6-bit command field can provide scope for expansion if needed. Multiple concurrent reductions may operate on the same tree. Each reduction can use a different reduction ID in the DFA. The wait count can be set to be sufficient for one contribution per node in a maximally sized system. In one embodiment, the wait count can be a 20-bits value.

[0082] In addition, four configurable timeout values can be provided. These 24-bit values can be given in units of 1024 clock cycles. At 850MHz, this provides a range from 1.20us to 20.2s. In addition, a 10-bit sequence number can be used, which can avoid being reused while an old reduction frame with the same multicast address and the same cookie value is still in the network. Note that this problem does not occur on the same reduction tree, because in-order

delivery prevents an old reduction frame from being delivered after frames transmitted later on the same tree. Therefore, for a late-arriving frame to cause a problem, the process it belongs to needs to exit first so that a newly launched process could request a new multicast tree, which can be built using the same multicast address. The chance that the old reduction frame survives the reprogramming of the multicast table is low. If it did survive, it still needs to intersect the new tree after the new tree performs 1024 reductions. Note that the 32-bit cookie value can also provide an extra layer of protection.

[0083] FIG. 7 shows an exemplary list of reduction operations, which are explained below. The FLT\_REP SUM and the MINMAXLOC operations support one reduction at a time per tree. For all other operations, four reductions per tree can be performed in parallel. The floating point sum operations can have four rounding modes and flush-to-zero (FTZ), which can be encoded in three bits as eight different commands. Floating point MIN and MAX have two modes for handling NaNs (mirroring the ARM MIN and MINNUM operations). FTZ only applies to denormalized results. If a result is to be denormalized, it is set to 0 instead and *flt\_inexact* is raised. The operation sometimes known as DAZ may be supported in software at the leaf nodes.

[0084] BARRIER: The data returned by BARRIER operations is always 0.

[0085] MINMAXLOC: The MINMAXLOC operators are used to support the MINLOC and MAXLOC. Operands 0 and 1 compute MINLOC, and operands 2 and 3 compute MAXLOC. FIG. 8 shows a set of MINMAXLOC operands. Note that when more than one index contains the minimum/maximum value, the lowest such index is recorded in the MINLOC/MAXLOC field.

[0086] FLT MIN and FLT MAX: When all inputs of FLT\_MIN or FLT\_MAX are floating point numbers, the minimum or maximum value is returned. When any operand of FLT\_MIN or FLT\_MAX is a NaN, NaN is returned. In a given pairwise reduction, if one operand is a signaling NaN and one operand is a quiet NaN, the signaling NaN is selected to be returned. The NaN returned can be turned into a quiet NaN with its sign bit cleared.

[0087] FLT MINNUM and FLT MAXNUM: These operations are similar to FLT\_MIN and FLT\_MAX but handle operands that are NaNs differently. In the absence of a signaling NaN, FLT\_MINNUM and FLT\_MAXNUM can return the smallest/largest numbers in the reduction. A quiet NaN is only returned when all of the operands in one of these reductions are quiet NaNs.

[0088] The behavior of signaling NaNs with regard to these operators can be controlled by a R\_TF\_RED\_CFG\_MODE register. In standard IEEE mode, a signaling NaN (SNaN) as an operand of a pairwise reduction always produces a quiet NaN as a result. This produces

indeterminate results for the complete reduction. In the recommended associative mode, when one operand is SNaN and the other is a number, the result is the number. Thus, in associative mode, if at least one operand in the reduction is a floating point number, the minimum or maximum floating point operand is returned. In either mode, if any operand is a signaling NaN, `flt_invalid` is returned.

**[0089]** FLT\_MINMAXNUMLOC: This operation computes both FLT\_MINNUM and FLT\_MAXNUM.

**[0090]** FLT\_SUM: This floating point sum operation has a flush-to-zero option and four rounding modes. When flush-to-zero is enabled, if the sum is denormalized, it is set to 0. The sign of the denormalized result is preserved. The four rounding modes match the ARM rounding modes and are shown in FIG. 9.

**[0091]** FLT\_REPSUM: The reproducible floating point sum is accomplished by splitting each floating point operand into up to four integer components, each of which has limited precision. The number of significant bits ( $W$ ) in each component is selected such that integer overflow cannot occur. The value of  $W$  is selected by software and is not observable in the hardware.  $W$  is used to compute the integer values  $IX$  to load into the reduction operand *rt\_data*. When the reduction is complete,  $W$  is used to construct the floating point result. The software may choose to set  $W$  to 40; in this case, up to  $2^{24}$  operands may be reduced. A floating point number can be represented by up to four  $W$ -bit integers as follows:  $\sum_{j=M}^{M+4} IX[j - M] \times 2^{W \times j}$ .

**[0092]** For each floating point operand, the software chooses the largest value of  $M$  such that the least significant bit of the operand appears in  $IX[0]$ . Software is responsible for loading the four  $IX$  values into *rt\_data* and the value of  $M$  in *rt\_repsum\_m*, which is an eight-bit signed integer.  $IX[0]$ , the least significant operand, is loaded into *rt\_data[0]*.

**[0093]** When two operands in this format are added by the reduction engine, if one operand's  $M$ ,  $M'$ , is larger than the other, the hardware can discard any  $IX[j]$  in the smaller operand where  $j < M'$  because these values may not have significance in the final result. If this occurs during the course of the reduction and any nonzero operands are dropped, *repsum\_inexact* can be returned.

**[0094]** When the reduction is complete, the root process can convert the resulting operands and *rt\_repsum\_m* into a floating point number. If there are more operands than are supported by the chosen  $W$ , *int\_overflow* may be returned. In this case, the result is not valid. Note that *int\_overflow* is only reported if the overflow occurs in one of the significant values returned in the result. The *rt\_repsum\_oflow\_id* identifies the most significant operand to overflow. When a partial result with *int\_overflow* is reduced with another partial result, if  $(M +$

rt\_repsum\_oflow\_id) of the overflow partial result is less than  $M'$  of the other partial result, the int\_overflow result code is dropped.

[0095] In some embodiments, static state for reduction operations can be programmed in the multicast table. This state can vary between devices in the same reduction tree. The state can be created by the management agent when a job starts or a new reduction tree is created. It can use the same mechanism as is used for setting up standard multicast entries. The multicast reduction trees are distinguished from multicast trees by a non-zero wait count. The timeout value for a reduction in a particular reduction engine can be determined by its wait count value and the configuration.

[0096] The reduction frame format is the same for *reduction\_arm* and *reduction\_data* frames. They can be distinguished by the *rt\_arm* field which can be set on all frames descending the tree. The 84-byte Portals-formatted reduction frame is shown in FIG. 10.

[0097] The Portals header and command fields are stored once in the reduction state. The 12-byte reduction header is shown in FIG. 11. The endianness of the operands, which is used for the MINMAXLOC reproducible sum operators, is shown in FIG. 12.

[0098] Errors or inexact results encountered during the course of a reduction are reported in *rt\_rc*. In general, these events do not prevent the reduction from completing. However, in the event of an opcode mismatch, the reduction cannot be performed. In this case, the Source Fabric Addresses (SFA) of two operands with differing opcodes are returned in *rt\_data*[0]. The reduction result codes are summarized in FIG. 13.

[0099] There is only one result code shared by the four parallel reductions. Result codes are defined in priority order. If a reduction encounters more than one exception condition, the largest is retained. For example, *flt\_invalid* is the highest priority FLT\_SUM result code.

[00100] In some embodiments, the reduction engine can be utilized from an Ethernet NIC which is generally not equipped to operate with the HPC Ethernet specification and utilize the Portals packet format. To facilitate reduction using reduction engines, the system can use the Soft Portals encapsulation, in which the Portals packet can be prepended with an Ethernet header constructed to be consistent with the configuration for Portals on the port that the NIC is connected to. This is illustrated in FIG. 14.

[00101] Within the Linux operating system, it is possible to construct these packets using a raw socket; this is suitable for functional testing since it requires the process to execute as root or with CAP\_NET\_RAW. The socket can be opened to receive only the specified ether type that a switch is configured to use in the Ethernet header it prepends to the Portals packets. It should be noted that the VNI field of the Portals packet is the protection mechanism, this can be

inserted in a privileged domain and for production usage this can be performed in a kernel module.

[00102] FIG. 15 shows an exemplary switching system that facilitates a reduction engine. In this example, a switch 1502 can include a number of communication ports, such as port 1520. Each port can include a transmitter and a receiver. Switch 1502 can also include a processor 1504, a storage device 1506, and a switching logic block 1508. Switching logic block 1508 can be coupled to all the communication ports and can further include a crossbar switch 1510 and a reduction engine logic block 1514.

[00103] Crossbar switch 1510 can include one or more crossbar switch chips, which can be configured to forward data packets and control packets among the communication ports. Reduction engine logic block 1514 can be configured to perform various dynamic reduction functions as described above. Also included in switching logic block 1508 is a multicast table 1516, which can store the reduction tree topology and state information to facilitate the reduction operations performed by reduction engine logic block 1514. Other types of data structure can also be used to store the topology and state information.

[00104] In summary, the present disclosure describes a switch capable of dynamically allocating a reduction engine in a network. The switch is equipped with a reduction engine that can be dynamically allocated based on a multicast frame to perform on-the-fly reduction. As a result, the network can facilitate an efficient and scalable environment for high performance computing.

[00105] The methods and processes described above can be performed by hardware logic blocks, modules, logic blocks, or apparatus. The hardware logic blocks, modules, logic blocks, or apparatus can include, but are not limited to, application-specific integrated circuit (ASIC) chips, field-programmable gate arrays (FPGAs), dedicated or shared processors that execute a piece of code at a particular time, and other programmable-logic devices now known or later developed. When the hardware logic blocks, modules, or apparatus are activated, they perform the methods and processes included within them.

[00106] The methods and processes described herein can also be embodied as code or data, which can be stored in a storage device or computer-readable storage medium. When a processor reads and executes the stored code or data, the processor can perform these methods and processes.

[00107] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the

above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.

**What Is Claimed Is:**

1. A switch, comprising:

a number of ports; and

an reduction engine coupled to at least one port and to:

5 be armed for a reduction process based on a received frame, wherein the received frame comprises a multicast address;

determine a number of output ports based on the multicast address; and

send the frame to the determined output ports, thereby allowing additional reduction engines to be armed for the reduction process.

10 2. The switch of claim 1, wherein the frame uniquely identifies the reduction process.

3. The switch of claim 1, wherein the reduction engine is further to set a wait count  
15 for the reduction process based on a number of endpoints below the switch based on a multicast tree corresponding to the multicast address.

4. The switch of claim 3, wherein the reduction engine is further to:  
identify one or more received reduction frames associated with the reduction process;  
20 determine that a total number of contributions indicated by the received reduction frames is equal to the wait count; and  
perform a reduction operation based on the contributions to generate a reduction frame to be forwarded to a parent node.

25 5. The switch of claim 4, wherein the generated reduction frame comprises a count of total contributions based on the received reduction frames.

6. The switch of claim 1, wherein the reduction engine is further to record a parent node from which the frame is received, thereby facilitating subsequent forwarding of a reduction  
30 frame toward a root node.

7. The switch of claim 1, wherein while determining the output ports the reduction engine is to look up a multicast table based on the multicast address.

8. A method, comprising:

arming a reduction engine for a reduction process based on a received frame,  
wherein the received frame comprises a multicast address;

determining a number of output ports based on the multicast address; and

5 sending the frame to the determined output ports, thereby allowing additional  
reduction engines to be armed for the reduction process.

9. The method of claim 8, wherein the frame uniquely identifies the reduction  
process.

10. The method of claim 8, further comprising setting a wait count for the reduction  
process based on a number of endpoints below the switch based on a multicast tree  
corresponding to the multicast address.

11. The method of claim 8, further comprising:  
identifying one or more received reduction frames associated with the reduction process;  
determining that a total number of contributions indicated by the received reduction  
frames is equal to the wait count; and  
performing a reduction operation based on the contributions to generate a reduction frame  
20 to be forwarded to a parent node.

12. The method of claim 8, wherein the generated reduction frame comprises a count  
of total contributions based on the received reduction frames.

13. The method of claim 8, further comprising recording a parent node from which  
the frame is received, thereby facilitating subsequent forwarding of a reduction frame toward a  
root node.

14. The method of claim 8, wherein determining the output ports comprises looking  
30 up a multicast table based on the multicast address.

15. A network system, comprising:  
a number of interconnected switches, wherein a respective switch comprises:

a number of ports; and

an reduction engine coupled to at least one port and to:

be armed for a reduction process based on a received frame, wherein the received frame comprises a multicast address;

- 5                   determine a number of output ports based on the multicast address; and  
                  send the frame to the determined output ports, thereby allowing additional  
reduction engines to be armed for the reduction process.

10           16.     The network system of claim 15, wherein the frame uniquely identifies the  
reduction process.

                  17.     The network system of claim 15, wherein the reduction engine is further to set a  
wait count for the reduction process based on a number of endpoints below the switch based on a  
multicast tree corresponding to the multicast address.

15                   18.     The network system of claim 15, wherein the reduction engine is further to:  
                  identify one or more received reduction frames associated with the reduction process;  
                  determine that a total number of contributions indicated by the received reduction frames  
is equal to the wait count; and  
20                   perform a reduction operation based on the contributions to generate a reduction frame to  
be forwarded to a parent node.

                  19.     The network system of claim 15, wherein the generated reduction frame  
comprises a count of total contributions based on the received reduction frames.

25                   20.     The network system of claim 15, wherein the reduction engine is further to record  
a parent node from which the frame is received, thereby facilitating subsequent forwarding of a  
reduction frame toward a root node.

30                   21.     The network system of claim 15, wherein while determining the output ports the  
reduction engine is to look up a multicast table based on the multicast address.




FIG. 1

2/12





FIG. 2

3/12

**FIG. 3A**

4/12




FIG. 3B

5/12




FIG. 4

6/12




FIG. 5A

7/12




FIG. 5B




FIG. 6

9/12

| Opcode  | Mnemonic         | Description                                                   |
|---------|------------------|---------------------------------------------------------------|
| 00 0000 | BARRIER          | No-op reduction                                               |
| 00 0001 | BIT_AND          | Bitwise AND                                                   |
| 00 0010 | BIT_OR           | Bitwise OR                                                    |
| 00 0011 | BIT_XOR          | Bitwise XOR                                                   |
| 01 0000 | INT_MIN          | Integer minimum                                               |
| 01 0001 | INT_MAX          | Integer maximum                                               |
| 01 0010 | INT_MINMAXLOC    | Integer min max with indexes                                  |
| 01 0100 | INT_SUM          | Integer sum                                                   |
| 10 0000 | FLT_MIN          | Floating point minimum                                        |
| 10 0001 | FLT_MAX          | Floating point maximum                                        |
| 10 0010 | FLT_MINMAXLOC    | Floating point min max with indexes                           |
| 10 0100 | FLT_MINNUM       | Floating point minimum number                                 |
| 10 0101 | FLT_MAXNUM       | Floating point maximum number                                 |
| 10 0110 | FLT_MINMAXNUMLOC | Floating point min max number with indexes                    |
| 10 1frr | FLT_SUM          | Floating point sum<br>f = flush-to-zero<br>rr = Rounding mode |
| 11 0000 | FLT_REFSUM       | Reproducible floating point sum                               |

FIG. 7

| Operand | Field  | Description            |
|---------|--------|------------------------|
| 0       | MIN    | Minimum value          |
| 1       | MINLOC | Index of minimum value |
| 2       | MAX    | Maximum value          |
| 3       | MAXLOC | Index of maximum value |

FIG. 8

| Mode | Name      | Description                    |
|------|-----------|--------------------------------|
| 0    | RND_NEAR  | round to nearest               |
| 1    | RND_CEIL  | round toward infinity          |
| 2    | RND_FLOOR | round toward negative infinity |
| 3    | RND_CHOP  | round toward zero              |

FIG. 9

10/12

| Field           | Description      | Byte offset | Bytes |
|-----------------|------------------|-------------|-------|
| portals_header  | Portals header   | 0           | 20    |
| portals_command | Portals command  | 20          | 16    |
| red_header      | Reduction header | 36          | 12    |
| rt_data         | Operands         | 48          | 32    |
|                 | FCS              | 80          | 4     |

FIG. 10

| Field                 | Description                                                    | Bit Location | Required in multicast arm | Size (bits) |
|-----------------------|----------------------------------------------------------------|--------------|---------------------------|-------------|
| rt_cookie             | Cookie value                                                   | 95-64        | Yes                       | 32          |
|                       | Padding                                                        | 63-61        |                           | 3           |
| rt_repsum_overflow_id | For REPSUM, the most significant operand to raise int_overflow | 60-59        |                           | 2           |
| rt_repsum_m           | Reproducible sum M                                             | 58-51        |                           | 8           |
| rt_rc                 | Result code                                                    | 50-47        |                           | 4           |
| rt_resno              | Result number                                                  | 46-37        |                           | 10          |
| rt_count              | Number of contributions                                        | 36-17        |                           | 20          |
| rt_op                 | Reduction operation                                            | 16-11        |                           | 6           |
| rt_arm                | Multicast arm command                                          | 10           | Yes                       | 1           |
| rt_seqno              | Sequence number                                                | 9-0          | Yes                       | 10          |

FIG. 11

| Field      | Description | Byte Offset | Bytes |
|------------|-------------|-------------|-------|
| rt_data[0] | Operand 0   | 0           | 8     |
| rt_data[1] | Operand 1   | 8           | 8     |
| rt_data[2] | Operand 2   | 16          | 8     |
| rt_data[3] | Operand 3   | 24          | 8     |

FIG. 12

| Result code | Name           | Ops affected                | Description                                                       |
|-------------|----------------|-----------------------------|-------------------------------------------------------------------|
| 0           | none           | all                         | No error                                                          |
| 1           | flt_inexact    | FLT_SUM                     | result was rounded                                                |
| 2           | reserved       |                             |                                                                   |
| 3           | flt_overflow   | FLT_SUM                     | result too large to represent                                     |
| 4           | flt_invalid    | FLT_SUM, FLT_MIN*, FLT_MAX* | An operand was a signaling NaN or two infinities were subtracted. |
| 5           | repsum_inexact | FLT_REPSUM                  | result was rounded                                                |
| 6           | int_overflow   | INT_SUM, FLT_REPSUM         | integer overflow                                                  |
| 7           | contr_overflow | all                         | Contributions exceed rt_waitcount                                 |
| 8           | op_mismatch    | all                         | rt_op mismatch                                                    |
| 9-15        | reserved       |                             |                                                                   |

FIG. 13




FIG. 14




FIG. 15

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2020/024260****A. CLASSIFICATION OF SUBJECT MATTER****H04L 12/911(2013.01)i, H04L 12/931(2013.01)i, H04L 12/801(2013.01)i, H04L 12/927(2013.01)i, H04L 12/741(2013.01)i, H04L 29/12(2006.01)i**

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)

H04L 12/911; G06F 017/50; G06F 15/76; G06F 9/46; H04J 3/14; H04L 12/803; H04L 12/875; H04L 12/931; H04L 12/801; H04L 12/927; H04L 12/741; H04L 29/12

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) &amp; keywords: reduction process, reduction frame, multicast address, switch, node, armed

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                      | Relevant to claim No. |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Y         | US 2013-0290673 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION)<br>31 October 2013<br>paragraphs [0051]-[0052]; and figure 3           | 1-21                  |
| Y         | US 2015-0193262 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION)<br>09 July 2015<br>paragraphs [0071], [0083]-[0084]; and figures 7, 10 | 1-21                  |
| A         | US 2003-0110455 A1 (JASON RAYMOND BAUMGARTNER et al.) 12 June 2003<br>paragraphs [0029]-[0031]                                          | 1-21                  |
| A         | US 2016-0142318 A1 (CRAY UK LIMITED) 19 May 2016<br>paragraphs [0023]-[0025]; and figure 1                                              | 1-21                  |
| A         | US 2008-0013453 A1 (MUNG CHIANG et al.) 17 January 2008<br>paragraphs [0015]-[0034]                                                     | 1-21                  |



Further documents are listed in the continuation of Box C.



See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"D" document cited by the applicant in the international application

"E" earlier application or patent 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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

06 July 2020 (06.07.2020)

Date of mailing of the international search report

**07 July 2020 (07.07.2020)**

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

YANG JEONG ROK

Telephone No. +82-42-481-5709



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2020/024260**

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s)                                                                                                                                                                       | Publication<br>date                                                                                                                                    |
|-------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| US 2013-0290673 A1                        | 31/10/2013          | US 2011-0296139 A1<br>US 8489859 B2<br>US 8601237 B2                                                                                                                                             | 01/12/2011<br>16/07/2013<br>03/12/2013                                                                                                                 |
| US 2015-0193262 A1                        | 09/07/2015          | US 2015-0193270 A1<br>US 9495204 B2<br>US 9495205 B2                                                                                                                                             | 09/07/2015<br>15/11/2016<br>15/11/2016                                                                                                                 |
| US 2003-0110455 A1                        | 12/06/2003          | US 6698003 B2                                                                                                                                                                                    | 24/02/2004                                                                                                                                             |
| US 2016-0142318 A1                        | 19/05/2016          | CN 102084627 A<br>CN 102084627 B<br>EP 2316202 A1<br>EP 2316202 B1<br>EP 3484108 A1<br>GB 2461132 A<br>GB 2461132 B<br>US 2011-0090797 A1<br>US 8908529 B2<br>US 9729450 B2<br>WO 2009-156720 A1 | 01/06/2011<br>07/10/2015<br>04/05/2011<br>15/08/2018<br>15/05/2019<br>30/12/2009<br>13/02/2013<br>21/04/2011<br>09/12/2014<br>08/08/2017<br>30/12/2009 |
| US 2008-0013453 A1                        | 17/01/2008          | US 7693072 B2                                                                                                                                                                                    | 06/04/2010                                                                                                                                             |