

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



A standard linear barcode is located at the bottom of the page, spanning most of the width. It is used for document tracking and identification.

**(43) International Publication Date
28 July 2011 (28.07.2011)**

(10) International Publication Number
WO 2011/091323 A1

(51) International Patent Classification:
G06F 15/00 (2006.01) G06F 15/76 (2006.01)

(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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(21) International Application Number: PCT/US2011/022152

(22) International Filing Date: 21 January 2011 (21.01.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 61/297,139 21 January 2010 (21.01.2010) US

(71) **Applicant** (for all designated States except US): QST

HOLDINGS, LLC [US/US]; 1800 Embarcadero Road, Palo Alto, CA 94303 (US).

(72) Inventors; and
(75) Inventors/Applicants (*for US only*): **MASTER, Paul, L.** [US/US]; 796 Dartshire Way, Sunnyvale, CA 94087 (US). **FURTEK, Frederick** [US/US]; 2470 Sharon Oaks Drive, Menlo Park, CA 94025 (US).

(74) **Agent: TANG, Wayne, L.**; Nixon Peabody LLP, 300 S. Riverside Plaza, 16th Floor, Chicago, IL 60606 (US).

(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: A METHOD AND APPARATUS FOR A GENERAL-PURPOSE, MULTIPLE-CORE SYSTEM FOR IMPLEMENTING STREAM-BASED COMPUTATIONS




FIG. 1

(57) Abstract: A method and system of efficient use and programming of a multi-processing core device. The system includes a programming construct that is based on stream-domain code. A programmable core based computing device is disclosed. The computing device includes a plurality of processing cores coupled to each other. A memory stores stream-domain code including a stream defining a stream destination module and a stream source module. The stream source module places data values in the stream and the stream conveys data values from the stream source module to the stream destination module. A runtime system detects when the data values are available to the stream destination module and schedules the stream destination module for execution on one of the plurality of processing cores.

A METHOD AND APPARATUS FOR A GENERAL-PURPOSE, MULTIPLE-CORE SYSTEM FOR IMPLEMENTING STREAM-BASED COMPUTATIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to Provisional Application Serial No. 61/297,139 filed January 21, 2010. This application is related to U.S. patent application Ser. No. 09/815,122, filed on Mar. 22, 2001, now U.S. Patent No. 6,836,839 entitled “ADAPTIVE INTEGRATED CIRCUITRY WITH HETEROGENEOUS AND RECONFIGURABLE MATRICES OF DIVERSE AND ADAPTIVE COMPUTATIONAL UNITS HAVING FIXED, APPLICATION SPECIFIC COMPUTATIONAL ELEMENTS”; U.S. patent application Ser. No. 10/384,486, now U.S. Patent No. 7,325,123 entitled HIERARCHICAL INTERCONNECT FOR CONFIGURING SEPARATE INTERCONNECTS FOR EACH GROUP OF FIXED AND DIVERSE COMPUTATIONAL ELEMENTS”; and U.S. patent application Ser. No. 10/443,501, now U.S. Patent No. 7,609,297 entitled “HARDWARE TASK MANAGER.” All of these applications are hereby incorporated by reference.

TECHNICAL FIELD

[0002] This invention relates in general to programming multiple processor systems and more specifically to a hardware task manager that efficiently utilizes parallel programming constructs incorporating both streams and threads.

BACKGROUND

[0003] A common limitation to processing performance in a digital system is the efficiency and speed of transferring instruction, data and other information among different components and subsystems within the digital system. For example, the bus speed in a general-purpose Von Neumann architecture dictates how fast data can be transferred between the processor and memory and, as a result, places a limit on the computing performance (e.g., million instructions per second (MIPS), floating-point operations per second (FLOPS), etc.).

[0004] Other types of computer architecture design, such as multi-processor or parallel processor designs require complex communication, or interconnection, capabilities so that each of the different processors can communicate with other processors, with multiple memory devices, input/output (I/O) ports, etc. With today's complex processor system designs, the importance of an efficient and fast interconnection facility rises dramatically.

[0005] However, such facilities are difficult to design to optimize goals of speed, flexibility and simplicity of design.

[0006] Currently, parallel programming is based on threads as the central, organizing principle of computing. However, threads are seriously flawed as a computation model because they are wildly nondeterministic and rely on programming style to constrain that non-determinism to achieve deterministic aims. Test and verification become difficult in the presence of this wild non-determinism. One solution has been suggested by GPU (Graphics Processing Unit) vendors is to narrow the forms of parallelism expressible in the programming model. Their focus on data parallelism, however, ties the hands of programmers and prevents exploiting the full potential of multi-core processors.

[0007] Further, threads do not just run on a bank of identical cores. A modern computer (supercomputer, workstation, desktop and laptops) contains a bewildering array of different heterogeneous cores all requiring separate programming models to program. For example, a motherboard may have one to four main CPUs (central processing units e.g., Pentium Processor) each having on-die 1 to 6 CPU cores with an on-die or on-package GPU (Graphics Processing Unit – e.g. NVIDIA GPU) which itself contains 16 to 256 GPU cores along with several discrete video & audio encode & decode cores (for the encoding and decoding of a multiplicity of video standards – e.g. MPEG2, MPEG4, VC-1, H.264 etc.). Also on the motherboard are from 1 to 4 discrete high end GPUs each containing 16 to 1024 GPU cores along with several discrete high-end configurable (meaning the core can be selected to encode/deocode a variety of pre-existing standards) video/audio encode & decode cores (for the encoding and decoding of a multiplicity of video standards – e.g. MPEG2, MPEG4, VC-1, H.264 etc., at very high resolutions and with multiple channels of sound). Additional subsystems composed of processing cores are added to the motherboard in the form of communications cores (e.g. TCP/IP offload cores which themselves are typical built from one or more CPU cores and one or more packet processing cores. WiFi cores, Blue Tooth cores, WiMax cores, 3G cores, 4G cores which are from one or more CPU cores and one or more broadband/baseband processing cores).

[0008] Current high end of the spectrum devices such as supercomputers add an additional processor in the form of one to four FPGAs (field programmable gate array) per motherboard. Each FPGA is itself composed of hundreds of thousand to tens of millions of very simplistic CLB processing cores along with multiple hard IP or Soft IP CPU core and multiple DSP cores). Then these motherboards themselves are then replicated and

interconnected in the hundreds to thousands to produce a modern supercomputer. These systems (either the desktops/workstations/laptops and/or the supercomputers) and then interconnected via the Internet to provide national and global computing capabilities.

[0009] The complexity of “managing” and “programming” such a diverse series of cores is a severe problem. Most programmers do not even attempt this and just settle for programming just one CPU core ignoring the rest of the cores. There are a certain number of algorithms known in the industry as “embarrassingly parallel problems” (e.g. the Google Search algorithm for example is simple to spread across multiple CPUs due to the fact that there is very little to no interactivity across the parallel threads). Unfortunately the vast majority of problems do not have these characteristics, they require a high degree of interactivity and synchronization across the multiple threads.

[0010] It would therefore be desirable to incorporate multithreading, unrestricted parallelism and deterministic behavior such as in modern programming language streams. Streams date at least to the introduction of the C programming language in 1978, and have been incorporated into such languages as C++, Java, Visual Basic and F#. However, in these languages, streams are relegated to a rather narrow role of providing a framework for I/O and file access. It is therefore desirable to expand the role of streams in parallel programming to first-class objects, a status roughly comparable to that of variables.

SUMMARY

[0011] According to one example, a programmable core based computing device is disclosed. The computing device includes a plurality of processing cores coupled to each other. A memory stores stream-domain code including a stream defining a stream destination module and a stream source module. The stream source module places data values in the stream and the stream conveys data values from the stream source module to the stream destination module. A runtime system detects when the data values are available to the stream destination module and schedules the stream destination module for execution on one of the plurality of processing cores.

[0012] Additional aspects of the invention will be apparent to those of ordinary skill in the art in view of the detailed description of various embodiments, which is made with reference to the drawings, a brief description of which is provided below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] FIG. 1 shows an overview of an adaptable computing engine compatible with a disclosed stream based programming model;

[0014] FIG. 2 illustrates a block diagram of an adaptive computing machine compatible with the programming model;

[0015] FIG. 3 is a diagram of a network word in the network of the adaptive computing machine in FIG. 2;

[0016] FIG. 4 illustrates the node wrapper interface between heterogeneous nodes and the homogenous network in the ACE architecture in FIG. 1 or ACM architecture in FIG. 2;

[0017] FIG. 5 shows basic components of a hardware task manager used for the node wrapper in FIG. 4;

[0018] FIG. 6 shows a point-to-point channel used for streaming data in the ACM architecture in FIG. 2;

[0019] FIG. 7 shows a point-to-point network word used by the point-to-point channel in FIG. 6;

[0020] FIGs. 8A-8D are diagrams of modules in relation to nodes for different stream flows;

[0021] FIGs. 9A and 9B are graphical representations of the assignments of values to streams;

[0022] FIG. 10 is a graphical depiction of a five-tap FIR filter that may be modeled using modules and stream concepts;

[0023] FIGs. 11A-C are modules with various configurations of FIFO;

[0024] FIG. 12 is a flow diagram of threads used in the example programming language; and

[0025] FIG. 13A-B are diagrams of forms of the join operation of the example programming language.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**ADAPTIVE COMPUTING ENGINE AND ADAPTIVE COMPUTING MACHINE**

[0026] FIG. 1 is a block diagram illustrating an example of a multi-processor system using an example computational model. Apparatus 100, referred to herein as an adaptive computing engine (ACE) 100, is preferably embodied as an integrated circuit, or as a portion

of an integrated circuit having other, additional components. In the exemplary embodiment, and as discussed in greater detail below, the ACE 100 includes one or more reconfigurable matrices (or nodes) 150, such as matrices 150A through 150N as illustrated, and a matrix interconnection network 110. Also in the exemplary embodiment, and as discussed in detail below, one or more of the matrices 150, such as matrices 150A and 150B, are configured for functionality as a controller 120, while other matrices, such as matrices 150C and 150D, are configured for functionality as a memory 140. The various matrices 150 and matrix interconnection network 110 may also be implemented together as fractal subunits, which may be scaled from a few nodes to thousands of nodes.

[0027] In a preferred embodiment, the ACE 100 does not utilize traditional (and typically separate) data, DMA, random access, configuration and instruction busses for signaling and other transmission between and among the reconfigurable matrices 150, the controller 120, and the memory 140, or for other input/output ("I/O") functionality. Rather, data, control and configuration information are transmitted between and among these matrix 150 elements, utilizing the matrix interconnection network 110, which may be configured and reconfigured, in real-time, to provide any given connection between and among the reconfigurable matrices 150, including those matrices 150 configured as the controller 120 and the memory 140.

[0028] The matrices 150 configured to function as memory 140 may be implemented in any desired or exemplary way, utilizing computational elements (discussed below) of fixed memory elements, and may be included within the ACE 100 or incorporated within another IC or portion of an IC. In the exemplary embodiment, the memory 140 is included within the ACE 100, and preferably is comprised of computational elements which are low power consumption random access memory (RAM), but also may be comprised of computational elements of any other form of memory, such as flash, DRAM, SRAM, MRAM, ROM, EPROM or E2PROM. In the exemplary embodiment, the memory 140 preferably includes direct memory access (DMA) engines, not separately illustrated.

[0029] The controller 120 is preferably implemented, using matrices 150A and 150B configured as adaptive finite state machines (FSMs), as a reduced instruction set ("RISC") processor, controller or other device or IC capable of performing the two types of functionality discussed below. Alternatively, these functions may be implemented utilizing a conventional RISC or other processor. The first control functionality, referred to as "kernel" control, is illustrated as kernel controller ("KARC") of matrix 150A, and the second control

functionality, referred to as "matrix" control, is illustrated as matrix controller ("MARC") of matrix 150B. The kernel and matrix control functions of the controller 120 are explained in greater detail below, with reference to the configurability and reconfigurability of the various matrices 150, and with reference to the exemplary form of combined data, configuration and control information referred to herein as a "silverware" module.

[0030] The matrix interconnection network 110 of FIG. 1, includes subset interconnection networks (not shown). These can include a boolean interconnection network, data interconnection network, and other networks or interconnection schemes collectively and generally referred to herein as "interconnect," "interconnection(s)" or "interconnection network(s)," or "networks," and may be implemented generally as known in the art, such as utilizing FPGA interconnection networks or switching fabrics, albeit in a considerably more varied fashion. In the exemplary embodiment, the various interconnection networks are implemented as described, for example, in U.S. Pat. No. 5,218,240, U.S. Pat. No. 5,336,950, U.S. Pat. No. 5,245,227, and U.S. Pat. No. 5,144,166. These various interconnection networks provide selectable (or switchable) connections between and among the controller 120, the memory 140, the various matrices 150, and the computational units (or "nodes") and computational elements, providing the physical basis for the configuration and reconfiguration referred to herein, in response to and under the control of configuration signaling generally referred to herein as "configuration information." In addition, the various interconnection networks (110, 210, 240 and 220) provide selectable or switchable data, input, output, control and configuration paths, between and among the controller 120, the memory 140, the various matrices 150, and the computational units, components and elements, in lieu of any form of traditional or separate input/output busses, data busses, DMA, RAM, configuration and instruction busses.

[0031] It should be pointed out, however, that while any given switching or selecting operation of, or within, the various interconnection networks may be implemented as known in the art, the design and layout of the various interconnection networks, in accordance with the present invention, are new and novel, as discussed in greater detail below. For example, varying levels of interconnection are provided to correspond to the varying levels of the matrices, computational units, and elements. At the matrix 150 level, in comparison with the prior art FPGA interconnect, the matrix interconnection network 110 is considerably more limited and less "rich," with lesser connection capability in a given area, to reduce capacitance and increase speed of operation. Within a particular matrix or computational unit,

however, the interconnection network may be considerably more dense and rich, to provide greater adaptation and reconfiguration capability within a narrow or close locality of reference.

[0032] The various matrices or nodes 150 are reconfigurable and heterogeneous, namely, in general, and depending upon the desired configuration: reconfigurable matrix 150A is generally different from reconfigurable matrices 150B through 150N; reconfigurable matrix 150B is generally different from reconfigurable matrices 150A and 150C through 150N; reconfigurable matrix 150C is generally different from reconfigurable matrices 150A, 150B and 150D through 150N, and so on. The various reconfigurable matrices 150 each generally contain a different or varied mix of adaptive and reconfigurable nodes, or computational units; the nodes, in turn, generally contain a different or varied mix of fixed, application specific computational components and elements that may be adaptively connected, configured and reconfigured in various ways to perform varied functions, through the various interconnection networks. In addition to varied internal configurations and reconfigurations, the various matrices 150 may be connected, configured and reconfigured at a higher level, with respect to each of the other matrices 150, through the matrix interconnection network 110. Details of the ACE architecture can be found in the related patent applications, referenced above.

[0033] Another example of an adaptive computing machine 160 that may use the parallel computational model is shown in FIG. 2. The adaptive computing machine 160 in this example has thirty-two heterogeneous leaf nodes 180 that are coupled together via a network 162. The network 162 has a single root 164 that is coupled to a group of network input ports 166, a group of network output ports 168, an optional system interface port 170, an external memory interface 172 and an internal memory interface 174. A supervisor node or K-node 178 is also coupled to the root 164.

[0034] The nodes 180 are each grouped in a quadtrees such as the quadtree 182. The quadtrees such as the quadtree 182 are implemented using 5-ported switch elements 184, each connected to a single parent and up to four children nodes 180. The switch elements implement a fair, round-robin arbitration scheme and provide pipelining with multi-level look-ahead for enhanced performance. In this example, the width of all paths is constant (51 bits), but the option is available to widen pathways as a tree is ascended, in the style of Leiserson's fat trees, in order to increase network bandwidth.

[0035] In this example all traffic on the network 162 is in the form of 51-bit network words as shown in the network word 188 shown in FIG. 3. The network word 188 has a route field 190, a security bit 192, a service field 194, an auxiliary field 196 and a payload field 198. The route field 190 is the destination address of the network word 188. The two high-order bits in the route field 190 are the chip ID. The security bit 192 allows peeks (reads) and pokes (writes) to configuration memory. The security bit 192 is set only for words sent by the K-Node 178. The service field 194 defines the type of service and the auxiliary field 196 is dependent on the service type. The service field 194 defines one of sixteen service types including Point-to-Point (PTP) which relates to streaming data and PTP acknowledgement which supports flow control for PTP data and causes a Consumer or Producer Count at the destination node to be incremented or decremented.

NODE WRAPPER

[0036] FIG. 4 illustrates the interface between heterogeneous nodes and the homogenous network in the ACE architecture in FIG. 1 or the ACM architecture in FIG. 2. This interface is referred to as a "node wrapper" since it is used to provide a common input and output mechanism for each node. A node's execution units and memory are interfaced with the network and with control software via the node wrapper to provide a uniform, consistent system-level programming model. In this example, the node 180 includes a memory 210 and an execution unit 212. Details of the node wrapper can be found in the related patent applications referenced, above.

[0037] In a preferred embodiment, each node wrapper includes a hardware task manager (HTM) 200. Node wrappers also include data distributor 202, optional direct memory access (DMA) engine 204 and data aggregator 206. The HTM coordinates execution, or use, of node processors and resources, respectively. The HTM does this by processing a task list and producing a ready-to-run queue. The HTM is configured and controlled by a specialized node referred to as a K-node 178 in FIG. 2 or control node (not shown). However, other HTM control approaches may be used.

[0038] The node wrapper in FIG. 4 makes the node 180 identical in outward appearance to all other nodes in the adaptive computing machine 160 in FIG. 2 or the adaptive computing engine 100 in FIG. 1 regardless of its internal structure or functionality. The node wrapper also relieves the execution unit 212 from having to deal with myriad activities associated with task management and network interactions. Among other things, the node wrapper is

responsible for disposing of each incoming network word such as the network word 188 in FIG. 2 in an appropriate fashion in one clock cycle.

[0039] The execution unit 212 in FIG. 4 is responsible for executing tasks (a task is equivalent to a module instance). The execution unit 212 may include a digital signal processor (DSP), a reduced-instruction-set (RISC) processor, a domain-specific processor, an application-specific integrated circuit (ASIC) or a reconfigurable (FPGA) fabric. Regardless of its form, the execution unit 212 interacts with the node wrapper through a standard interface.

[0040] The nodal memory 210 is accessible to both the node wrapper and the execution unit 212. The nodal memory 210 is where the node wrapper deposits incoming streaming data and where the execution unit 212 accesses that data. A node's own memory 210, however, is typically not where the execution unit 212 sends output data. To minimize memory accesses, output data is usually sent directly to the node(s) requiring that data: the consumer node(s). Nodal memory 210 is also used to store task parameters and is available to tasks for temporary (scratchpad) storage.

[0041] In a multi-node system such as the ACM 160 in FIG. 2 where the nodes 180 are both consumers and producers of streaming data, matching production and consumption rates is desirable. A producer task on one node may produce data at a rate that is either greater than or less than the rate at which a consuming task on another node can handle. If the producer is sending data at a greater rate than the consumer can handle, then data is eventually lost. If the producer is sending data at a lesser rate than the consumer can handle, then the consumer may be starved for data, thereby potentially causing the consumer to sit idle waiting for additional data.

[0042] The ACM 160 provides, via the Point-to-Point protocol, and the node wrapper in FIG. 4 a single, uniform and consistent mechanism for task management, flow control and load balancing. Task management ensures that a task is placed in execution only when it has sufficient input data and when there is sufficient space in the consumer node(s) to accommodate the data produced by the task. Flow control guarantees that a producer task will never overwhelm a consumer task with too much data in too short a time. Load balancing permits a producer task to distribute data among several alternate consumer nodes, thus allowing the producer task to operate at a potentially higher rate.

[0043] Streaming data is transferred between two nodes 180 (points) via a point-to-point channel (point-to-point stream) 250 as shown in FIG. 5. Each PTP channel such as the

point-to-point channel includes a producer node 252, a producer task 254, an output port 256, a consumer node 258, an input port 260, an input buffer 262 and a consumer task 264. The producer task 254 runs on the execution unit of the producer node 252 and produces a finite-sized block of PTP data per task activation. The block of data is sent over the PTP channel 250 as a sequence of PTP words. The sending of blocks is shown as task 1 in FIG. 5. The output port 256 on the producer node 252 is associated with the producer task 254.

[0044] The consumer task 264 receives PTP data from the PTP channel 252 via the input port on the consumer node 258. The circular input buffer 262 in the nodal memory of the consumer node 258 stores the incoming PTP data. A consumer task such as the consumer task 264 runs on the execution unit of the consumer node 258 and consumes a finite amount of the PTP data residing in the circular input buffer 262 per task activation (Task 2 in FIG. 5).

[0045] Data is conveyed over the PTP channel 252 when the producer task 254 transfers a 50-bit point-to-point word 270 as shown in FIG. 6 to the node wrapper in the producer node 252. The point-to-point word 270 has the same fields as the network word 188 in FIG. 3 and like elements/fields are labeled with the same element numbers in FIG. 5. The point-to-point word 270 includes a node word 272 in the route field 190, a port word 274 in the auxiliary field 196 and a data word 276 in the payload field 198. In this example, a 51st bit, the security bit 192, is added later by the network 162 in FIG. 2. The node wrapper such as the node wrapper in FIG. 4, in turn, hands the PTP word over to the packet-switched network for transfer to the consumer node 258 in FIG. 5. The 8-bit route field 190 of the PTP word 270 provides the address in the form of the node word 272 of the consumer node such as the node 258 in FIG. 5. The port word 274 includes the low-order 5 bits of the auxiliary field 196 that indicate to which of the input ports of the consumer node the data is directed. When the PTP word arrives at the consumer node, the node wrapper deposits the 32-bit data word 176 in the payload field 198 into the circular input buffer associated with the indicated input port. The transfer is then complete.

[0046] The ACM 160 includes mechanisms for task management, flow control and load balancing. There is an input buffer associated with each input port. There is also a two's-complement signed count associated with each port, both input and output.

[0047] For an input port, the count is referred to as a consumer count since it reflects the amount of data in that port's input buffer that is available to be consumed by the associated task. A consumer count is enabled when its value is non-negative – that is, when its sign bit is 0. An enabled consumer count indicates that the associated input buffer has the minimum

amount of data required by an activation of the associated task. At system initialization, or upon reconfiguration, a consumer count is typically reset to $-C$, where C is the minimum number of 32-bit words required per task activation.

[0048] For an output port, the count is referred to as a producer count since it reflects the amount of available space in the downstream input buffer to accept the data that is produced by the associated task. A producer count is enabled when its value is negative – that is, when its sign bit is 1. An enabled producer count indicates that the downstream input buffer has space available to accommodate the maximum amount of data produced per activation of the associated task. At system initialization, or upon reconfiguration, a producer count is typically reset to $P - S - 1$, where P is the maximum number of 32-bit words produced per task activation and S is the size of the downstream input buffer in 32-bit words.

[0049] Both consumer counts and producer counts are typically initialized to negative values, causing the consumer counts start out disabled while producer counts start out enabled. This initial state reflects the fact that input buffers are usually empty at system initialization/reconfiguration.

[0050] Consumer and producer counts are updated by a system of credits and debits in the form of forward acknowledgements and backward acknowledgements. Both types of acknowledgements are network words such as the acknowledgment network word 280 shown in FIG. 7. The acknowledgment network word 280 has the same fields as the network word 188 in FIG. 3 and like elements/fields are labeled with the same element numbers. The acknowledgment network 280 word is sent by a task as the last steps in a task activation. In both cases, the payload field 198 contains four subfields: a type of acknowledgement subfield 282 (one bit), a port subfield 284, (3) a task subfield 286 and an Ack Value subfield 288.

[0051] The sequence of acknowledgements that a task performs at the end of each activation is described below. For each output port of the task, a forward acknowledgement is sent to the consumer node specifying the consumer input port and the consumer task. The Ack Value is the number of PTP words the task just sent to the consumer input port. A backward acknowledgement (a self ack) is sent to the node on which the task resides specifying the output port and the task. The Ack Value is the number of PTP words the task just sent via the output port.

[0052] For each input port of the task, a backward acknowledgement is sent to the producer node specifying the producer output port and producer task. The Ack Value is minus the number of 32-bit words the task just consumed from the input port's buffer. A

forward acknowledgement (a self ack) is sent to the node on which the task resides indicating the input port and the task. The Ack Value is minus the number of 32-bit words the task just consumed from the input port's buffer.

HARDWARE TASK MANAGER

[0053] The hardware task manager 200 shown in FIG. 4 is the part of the node wrapper responsible for updating consumer and producer counts in response to incoming acknowledgements. It also monitors the sign bits of those counts and launches a task when an appropriate set of counts is enabled. This last responsibility is met using two signed counts that are associated not with a port but with a task: a task input count and a task output count. A task's input (output) count reflects the number of task consumer (producer) counts that are enabled. A task count is enabled when its value is non-negative. A task is enabled, and available for execution, when both its input count and its output count are enabled.

[0054] Incoming acknowledgements update various counts and cause tasks to be launched as follows. If a forward acknowledgement is received, the specified port is interpreted as an input port, and Ack Value is added to the corresponding consumer count. If the consumer count makes a transition from disabled to enabled (enabled to disabled), then the input count of the specified task is incremented (decremented) by 1. If a backward acknowledgement is received, the specified port is interpreted as an output port, and the Ack Value is added to the corresponding producer count. If the producer count makes a transition from disabled to enabled (enabled to disabled), then the output count of the specified task is incremented (decremented) by 1. If after a forward or backward acknowledgement is received, the specified task's input and output counts are both enabled, then the task is placed on the ready-to-run queue if it is not already on the queue. The task is launched when it reaches the head of the queue.

[0055] These actions embody the firing rule for tasks. They cause a task to be placed on the ready-to-run queue and ultimately executed when a sufficient number of consumer counts and a sufficient number of producer counts are enabled. What those sufficient numbers are is determined by the initial values of a task's input count and output count. If I (O) is the number of input (output) ports associated with a task and $IC_{Initial}$ ($OC_{Initial}$) is the initial value of the task's input (output) count, and if it is assumed all consumer counts are initially disabled and all producer counts are initially enabled as discussed above, then a task fires when

$-IC_{Initial}$ out of I consumer counts are enabled

AND

$(O - OC_{Initial})$ out of O producer counts are enabled

For example, for I = 4,

If $IC_{Initial} = -1$, then 1 out of 4 consumer counts must be enabled

If $IC_{Initial} = -2$, then 2 out of 4 consumer counts must be enabled

If $IC_{Initial} = -3$, then 3 out of 4 consumer counts must be enabled

If $IC_{Initial} = -4$, then 4 out of 4 consumer counts must be enabled

For O = 4,

If $OC_{Initial} = 3$, then 1 out of 4 producer counts must be enabled

If $OC_{Initial} = 2$, then 2 out of 4 producer counts must be enabled

If $OC_{Initial} = 1$, then 3 out of 4 producer counts must be enabled

If $OC_{Initial} = 0$, then 4 out of 4 producer counts must be enabled

[0056] The programming of the multi-processor system such as the ACE 100 in FIG. 1 and the ACM 200 in FIG. 2 may be accomplished using what may be termed as the Stream C programming language.

STREAM C MODULES

[0057] In a Stream C program, there is only one mechanism for expressing concurrency: through the concurrent operation of the program's modules (and module-like stream expressions). Syntactically, modules are very similar to C functions, but semantically, they are different. A C function (subroutine) initiates activity only when it is called. In a call, control is passed to the C function, usually together with some input arguments. The C function then performs a task/computation, and when finished, returns control together with any output result. Unlike C functions, modules are not called nor is control passed to or returned from modules. Instead, modules carry on ongoing interactions with other modules and the outside world through their input and output ports. Through these ports, a module receives streams of input values and emits streams of output values.

[0058] The syntax of a module prototype is identical to that of a C function prototype, with three exceptions. First, the keyword, `stream`, precedes a module prototype. This tells the compiler/linker that each module input and module output is associated with a stream of values of the specified type, not an individual value. Second, to permit a module to have multiple output streams, the module return type may be replaced by a parentheses-delimited

list that is identical in syntax to the input parameter list. Third, to extend the notion of an array to modules, a square-bracket-delimited list of array indices may be inserted immediately after the module name and before the input argument list. The module arrays are discussed below.

[0059] The following are two examples of module declarations:

```
stream int moduleA(int, int);

stream (int, int) moduleB(int, int);
```

Parameter names have been omitted here since they are not required in a module declaration (in contrast to a module definition or module instantiation), but Parameter names may be included at the programmer's discretion, usually as a mnemonic aid, for inputs and, when there are multiple outputs, for the outputs as well. The two declarations, for example, might then be expressed as:

```
stream int moduleA(int a, int b);

stream (int x, int y) moduleB(int a, int b);
```

[0060] The first declaration indicates that moduleA has two input streams, both of integer type, and a single output stream, also of integer type. The second declaration indicates that moduleB has two input streams, both of integer type, and two output streams, also both of integer type.

[0061] Like the definition of a C function, the definition of a module has a body delimited by curly braces ({ and }). Also as in the definition of a C function, each module input (and output in the case of multiple outputs) is assigned an identifier. The following are two examples of module definitions:

```
stream int moduleA(int a, int b)
{
    // Module Body

}

stream (int x, int y) moduleB(int a, int b)
{
    // Module Body
```

}

[0062] A module instantiation is the module counterpart to a C function call. Like a function call, a module instantiation is where a module gets used. While the syntax of these two types of expressions are similar, the semantics are different. A fragment of C code may be expressed as:

```
int x, y;
int F(int, int);
.
.
.
int z = F(4, x + 5*y);
```

[0063] The first statement declares that x and y are integers, while the second declares that F is a function with two integer parameters and an integer result. The last statement is an assignment containing the function call $F(4, x + 5*y)$, which has two arguments, the expressions 4 and $x + 5*y$, corresponding to the two parameters of F .

[0064] The stream version of this code fragment is as follows:

```
stream int x, y;
stream int F(int, int);
.
.
.
stream int z = F(4, x + 5*y);
```

[0065] In the stream version, each of the statements above is prefaced with the keyword `stream`. The change in syntax produces a dramatic change in semantics. Instead of individual values, streams of values are used. Thus the first statement declares that x and y are integer streams, while the second declares that F is a module with two integer-stream inputs and an integer-stream output. The last statement is now an assignment containing the module instantiation $F(4, x + 5*y)$, which has two arguments, the stream expressions 4 and $x + 5*y$, corresponding to the two parameters of F .

[0066] In the case of the function call, each execution of the assignment $z = F(4, x + 5*y)$ causes expressions 4 and $x + 5*y$ to be evaluated and the two resulting values to be supplied as parameters in a call to function F . After some period of time, F returns a value. In the case of the module instantiation, there is no execution of the assignment $z = F(4, x +$

$5*y$) and no call to module F. Instead, an instance of module F is created (instantiated) at system initialization, just before the Stream C program begins execution, thereby making the instance ready to receive streams of integers on its two input ports and produce a stream of integers on its output port. And once program execution commences, the instance of F remains operative until program termination i.e., the instance of F is persistent.

[0067] This simple example illustrates the general mechanism used in Stream C to create a community of interacting modules. Each module instantiation causes a separate module instance to be created at system initialization. Once created (instantiated), a module instance is ready to receive streams of values on its input ports and produce streams of values on its output ports. Furthermore, once program execution commences, the module instance remains operative until program termination.

[0068] The general form of the instantiation of a module with multiple output ports is:

(<identifier-list>) <module-identifier> (<expression-list>)

While the input arguments are expressions, the output arguments are identifiers. These identifiers serve to name the otherwise nameless output streams. The stream assignment statement above plays the same role by assigning the name z to the otherwise nameless output stream of F (4, x + 5*y). For example:

```
stream int w, x, y, z;
stream (int, int) F(int, int);
.
.
.
(w, z) = F(4, x + 5*y);
```

[0069] As before, F has two integer-stream inputs, but in contrast to the earlier example, F now has two integer-stream outputs. Those two output streams appear in the instantiation of F as the list of identifiers (w, z), which serves to give the two output streams the names w and z.

[0070] Statements within a module body fall into two categories (or domains), stream statements that involve only streams and thread statements that include the entire range of C statements as well as statements that allow a thread to read from and write to streams. Because each module instantiation causes a separate module instance to be created at system initialization, Stream C does not allow a module to have within its body, or within the body of a submodule, an instantiation of itself. In other words, circularity in module references is

not allowed. This prohibition helps avoid the difficult task of instantiating an infinite number of module instances.

[0071] In a Steam C module, there is no notion of returning control, and so the `return` statement is inappropriate. In a module, output values are simply inserted into a module output stream. But in order to do that, the output stream must have a name. For modules with a parentheses-delimited list of named output streams, that's not a problem. It is a problem, however, when the module prototype provides only the type of the module's sole output stream. In that case, code in the module's body, either in the stream domain or thread domain, can use the keyword `out` as the name of the default module output stream. This usage is illustrated in the following code fragment.

```
stream int moduleA(int a, int b)
{
    .
    .
    .
    out = a + b;
    .
    .
    .
}
```

[0072] Modules and the streams they are partnered with provide the framework for the web of interactions and concurrent activity typical of a Stream C program, while functions provide the computational building blocks of a program. Although modules deal with streams of values, that does not prevent modules from accessing individual values within a stream and supplying those values to a function. Likewise, a module can access the output value of a function and insert that value into a stream. A function, on the other hand, cannot reference a module because there is no mechanism within a function for such interaction. Because of this asymmetry in capabilities, modules are found at the higher levels of a program hierarchy, while functions are found at the lower levels.

[0073] While the differences between modules and functions are substantial, there is one area in which they are similar. They both support side effects, that is, they both may manipulate external data structures independently of their input and output ports. This stems from the fact that modules may contain threads that may have side effects.

[0074] FIG. 8A shows a generic module which includes a module 300 and a number of input streams 302 (from zero to N) which provide data/control to the module 300 and a number of output streams 304 (from zero to N) which provide data/control to the next

module/function. A module with no output streams is a “sink” and a module with no input streams is a “source.”

[0075] FIG. 8B is a diagram showing two modules, module A 300 and module B 310 each having corresponding input streams 302 and 312 and output streams 304 and 314. The output streams 304 of Module A 300 are attached to the input streams 312 of Module B 310. The Module A 300 is mapped to run on a CPU core 308 and the Module B 310 is mapped to run on a second CPU core 318. The cores 308, 318 and 328 are similar to the nodes 180 in FIG. 2.

[0076] FIG. 8C is a diagram that shows the mapping of both module A 300 and module B 310 onto the same CPU core such as the CPU core 308. In this instance the modules 300 and 310 act like any other separate thread of control. An operating system which is running on the second core 318 may schedule the modules 300 and 310 based on preemptive multi-tasking or run to completion/release. Since both modules 300 and 310 and input/output streams 302, 312 and 304, 314 are “persistent” (that is they remain ready to perform processing) additional information must be provided to a conventional operating system about when to schedule a module based on both the availability of “enough” input stream data to perform a computation and “enough” space so that the output stream can deliver the computed data.

[0077] A variety of different algorithms can be used to perform the Module to Core Mapping. These may include cache proximity where modules which share the greatest number of streams are placed in cores which share a L1 cache followed by a shared L2 cache followed by a shared L3 cache followed by a shared DRAM. They may also include a physical proximity algorithm where modules which share the greatest number of streams are placed in cores which are physically adjacent to each other. For example the algorithm may start with the die and then the integrated circuit on the motherboard, then the motherboard in the rack, then the rack in the floor of the building, then the building geographically adjacent. Another algorithm may be the next available free where modules are allocated to cores based on the next “free” core based on either CPU usage (current or average weighted over time) or the next sequentially available core. Another algorithm may be a predictive load that selects the modules and cores based on estimated statistical sampling. A running average of core utilization is used to load modules to the lightest loaded core. Another algorithm may be user specified where a user specified virtual core ID is used to place all modules onto a physical

core ID. When the number of virtual core ID's exceeds the physically available cores then multiple modules are evenly loaded across all available physical cores.

[0078] FIG. 8D show various data structures 330, 332 and 334 that may reside in the module A 300 and are used in the input and output streams 302 and 304. The data structures 330, 332 and 334, which reside either in memory/cache or a TLB, contain the critical information which is required for either a single core or a multi-core system to schedule and transport data from input streams such as the input stream 302 to output streams such as the output stream 304, input streams 302 into modules such as the module A 300, and the module A 300 out to output streams 304. For each module there is information that uniquely identifies the module, uniquely identifies all the input streams into that module, uniquely identifies all the output streams of that module, uniquely identifies how input and output streams are "connected," uniquely identifies the core, and maintains state information so that a module may be relocated from a core to another core, or swapped out via virtual memory. Streams can be dynamically added or deleted from modules, and modules can be dynamically added or deleted from cores.

STREAMS

[0079] The term stream in the Stream C programming language refers to a sequence of data values, all of the same type and typically made available over time. In Stream C, however, streams provide much more than a framework for input and output. Streams are elevated to first-class objects, a status roughly comparable to that of variables. This means that a stream may be bound to an identifier (i.e., it can be named), an input parameter of a function (i.e., an input parameter of a module), the output of a function (i.e., an input parameter of a module), a parameter within an expression and the output of an expression

[0080] A stream conveys values of a single type from one or more stream sources to one or more stream destinations. The precise details of how this transport is accomplished is implementation dependent, and depends upon, among other things, whether the stream is confined to a single semiconductor die or whether the stream spans several meters or possibly even thousands of kilometers. Except when dealing with performance issues, a programmer need not be concerned with those details, and need only be concerned with those aspects of a stream that relate to four stream attributes: the stream type, the stream name, the stream sources and the stream destinations.

[0081] The stream type indicates the type of values being conveyed. The type which may be any legitimate C type, including pointers and types defined via `typedef`, may be specified implicitly by context for example, by appearing as a module input or output parameter or explicitly using a stream declaration as described below.

[0082] The stream sources are the points at which values are placed into the stream. Possible stream sources are include an input parameter of a module definition, an output of a module instantiation, the output of a stream expression and a thread (discussed below). The stream destinations are the points to which the stream conveys values. Possible stream destinations include an output parameter of a module definition, an input argument of a module instantiation, an input of a stream expression and a thread. An optional stream name is a name/identifier is assigned to a stream when the stream appears as a module input or output parameter or when it is introduced in a stream declaration. An example of an unnamed stream is the output stream of a stream expression that has not been assigned a name via a stream assignment.

[0083] The notion of stream attributes is illustrated in the following code fragment containing a declaration of a function `F` and a partial definition of a module `M`.

```
stream int F(int, int);

stream (int zStrm) M(int xStrm, int yStrm)
{
    :
    :
    zStrm = xStrm*yStrm + F(xStrm, yStrm);
    :
    :
}.
```

[0084] There are three named streams here: `xStrm`, `yStrm` and `zStrm`, all of type `int`. `xStrm` and `yStrm` each have a single source: an input parameter of module `M`. The destinations of `xStrm` and `yStrm` are represented by the two instances of `xStrm` and `yStrm`, respectively, appearing in the assignment expression in the body of `M`. (Recall that, in C, an assignment is also an expression.) Those instances represent inputs to the assignment expression. `xStrm` and `yStrm` thus each have a single source and two destinations.

[0085] A stream expression is identical to a C expression, except that in place of variables, there are now input streams. A stream expression also has an output stream, which

carries the results from expression evaluations. By default, this output stream is nameless, but it can be assigned a name using a stream assignment, which is just what we've done in the above assignment. Thus the output stream of the stream expression

$$xStrm * yStrm + F(xStrm, yStrm)$$

is assigned the name `zStrm` by the stream assignment

$$zStrm = xStrm * yStrm + F(xStrm, yStrm)$$

Either of these two expressions may be considered the source of `zStrm`. The destination of `zStrm` is the output stream of module `M`, which is denoted by the output parameter `zStrm` of module `M`.

```
stream (int zStrm) M(int xStrm, int yStrm)
```

`xStrm` thus has a single source and a single destination.

[0086] The most crucial properties of a stream relate to the stream's role as a conveyor of values. There are four such properties: a) values do not enter a stream except at stream sources or at system initialization using the `initialize()` function; b) Values entering a stream at a single source are totally ordered in time; c) once entered into a stream, a value is eventually delivered to all stream destinations; if there are multiple destinations, then a separate copy of the value is delivered to each destination; and d) values from a single source are received at each stream destination in the same order in which they entered the stream i.e., there is no leapfrogging of values in a stream. These four properties are the only guarantees that a stream provides regarding the transport of values. Any other property that does not follow as a logical consequence from these four is not a general stream property.

[0087] Because a stream is obliged to only deliver values eventually, the latency of a stream, the time it takes a value to travel from a stream source to a stream destination, is indeterminate. In fact, the latency may vary with time and between different source-destination pairs of the same stream. Fixed or at least, bounded latencies, however, can still be achieved by relying upon guarantees provided by the system implementation (rather than the programming model). A source-destination pair confined to a single semiconductor die, for example, will usually have bounds on its latency.

[0088] The above four properties also have implications for stream determinacy and indeterminacy (non-determinism). For a stream with a single source, the four properties ensure deterministic stream behavior. That means that the order in which values are placed

into a single-source stream completely determines the order in which values are delivered to all stream destinations. For a stream with multiple sources, however, the situation is very different. To illustrate the issues arising from multiple stream sources, consider the following adaption of the code fragment from the preceding section. (`out` is the default output stream of a single-output module).

```
int F(int);
stream int M(int xStrm, int xStrm)
{
    :
    :
    :
    out = xStrm*xStrm + F(xStrm);
    :
    :
    :
}
```

[0089] The two input parameters of module M are the same: `xStrm`. From the four properties, values entering `xStrm` through the first input parameter of module M will be received at each of the three destinations of `xStrm` in the same order in which they entered the stream. Values entering `xStrm` through the second input parameter of module M will be received at each of the three destinations of `xStrm` in the same order in which they entered the stream. That means that the two streams of values are merged or interleaved before reaching each destination of `xStrm`.

[0090] How interleaving is carried out is influenced, in general, by program structure. The missing parts of the program above, for example, may be structured in a way that leads to an interleaving that strictly alternates between parameter-one and parameter-two values. So, for example, if the integers arriving on the two input parameters (streams) `xStrm` of module M begin with the sequences:

<u>xStrm Source One</u>	<u>xStrm Source Two</u>
6	3
8	5
1	9
:	:
:	:
:	:

then the sequence arriving at each of the three destinations of `xStrm` in the expression

`out = xStrm*xStrm + F(xStrm)`

might begin with

xStrm Destinations

6
3
8
5
1
9
.
.

Such program-imposed determinism though is not always the case, and there are situations in which values from multiple stream sources are interleaved non-deterministically. Moreover, depending on the target system, those nondeterministic interleavings may differ from one stream destination to another. Thus, for example, if the values arriving on the two input parameters (streams) of module M are the same as above, then the sequence arriving at the three destinations of xStrm might begin with:

<u>xStrm Destination One</u>	<u>xStrm Destination Two</u>	<u>xStrm Destination Three</u>
6	3	3
8	5	6
1	6	5
3	8	8
5	1	1
9	9	9
.	.	.
.	.	.
.	.	.

This indeterminacy in order of arrival of values at the destinations of a multi-source stream contrasts with the fixed order of arrival across all destinations of a single-source stream. That fixed arrival order allows adopting notation that is useful in below. For a single-source stream `ssStrm` and a non-negative integer `i`,

`ssStrm(i)`

denotes the i^{th} value appearing at all destinations of `ssStrm`. By convention, `ssStrm(0)` denotes the first value appearing at all destinations.

[0091] When a value arrives at a stream destination, if the destination is an output parameter of a module definition or an input argument of a module instantiation, then the value is handed over to a stream on the other side of a module boundary. The value thus remains in transit. If the destination is an input of a stream expression or a thread then the value comes to rest in a FIFO queue.

[0092] To illustrate remaining in transit, the following code fragment is shown.

```
stream int module1(int);

stream int module2(int xStrm)

{
    :
    :
    :

    out = module1(xStrm);

    :
    :
    :
}
```

The code fragment includes two modules, `module1` and `module2`, each with a single input stream and single output stream and two named streams, `xStrm` and `yStrm`, both within the definition (body) of `module2`. The sole destination of `xStrm: module1(xStrm)` is an input argument of an instantiation of `module1`. A value arriving at this destination is simply passed across `module1`'s boundary to an internal stream of `module1`. The situation is similar for values arriving at the sole destination of `yStrm`:

```
stream (int yStrm) module2(int xStrm)
```

Since this destination is an output parameter of a `module2`, arriving values are simply passed across `module2`'s boundary to a stream external to `module2`.

[0093] Another example is the case where a stream destination is an input of a stream expression such as the following code fragment.

```
stream int F(int, int);

stream int M(int xStrm, int yStrm)

{
    :
    :
    :

    out = xStrm*yStrm + F(xStrm, yStrm);

    :
    :
```

}

Within the body of module M is the stream expression

$$xStrm * yStrm + F(xStrm, yStrm)$$

which contains two destinations of $xStrm$ and two destinations of $yStrm$. It also contains the two operators $*$ and $+$ and the function F , which are ordinary C constructs. That means that in order to evaluate this expression, the two operators and the function F must be supplied with individual values.

[0094] The queues are automatically inserted by the Stream C linker/loader and are managed by the Stream C runtime. Among the responsibilities of the runtime is signaling when a queue is empty and ensuring that no queue ever overflows. Each queue is guaranteed to have a capacity of at least two values of the associated data type, although a programmer may request a specific amount via a pragma as described below. In the stream of this example, there are four queues, one for each of the four stream destinations (stream expression inputs). These queues are largely invisible to the programmer.

[0095] Once a Stream C program begins executing (operating), the only way for a value to enter a stream is through a stream source. One or more streams may form a directed cycle which requires a value already in the stream. The simplest such cycle occurs when a stream appears on both sides of a stream assignment as in:

$$xStrm += yStrm$$

which is equivalent to

$$xStrm = xStrm + yStrm$$

[0096] FIG. 9A is a first graphical representation 400 of this assignment where the directed cycle consists of the feedback path from the output of the $+$ operator to one of the two inputs of the same operator. It is the lack of a value on this path that prevents the $+$ operator from consuming a value from each input stream and producing a value on the output stream. So unless a value 404 is placed in the feedback path as shown in the second graphical representation 402 before execution begins, the $+$ operator will never fire.

[0097] Another issue relates to changing the offset of one single-source stream relative to another single-source stream. For example, if $aStrm$ and $bStrm$ are both inputs to the same module or stream expression, as in

$$aStrm + bStrm$$

which is represented graphically in FIG. 9B, and the module or expression consumes values from the streams in pairs, one from aStrm and one from bStrm. If it is desired that aStrm(n) (i.e., the n^{th} value arriving on aStrm) to match up with bStrm(n+2) (i.e., the $n+2^{\text{nd}}$ value arriving on bStrm). Thus aStrm(0) would match up with bStrm(2), aStrm(1) would match up with bStrm(3), and so on.

[0098] The solution to both issues is provided by the stream initialization statement, which has the form

```
<stream-identifier>.initialize(<value-list>);
```

[0099] When the Stream C compiler/linker/loader encounters this statement, it takes the statement as a directive to insert a FIFO queue at each destination of <stream-identifier> whether the destination is an output parameter of a module definition, an input argument of a module instantiation, an input of a stream expression or a thread; size the queue at each stream destination so that it is sufficient to hold at least $n+1$ values of type T, where n is the number of values in <value-list> and T is the type of <stream-identifier>; and place the values in <value-list> into the queue in order, with the first value in <value-list> placed at the front (head) of the queue.

[00100] For example, in FIG. 9A, to prevent deadlock, a value 404 is inserted in the feedback path and also in the expression output in the graphic representation 402 by initializing xStrm with the value 0 via the statement

```
xStrm.initialize(0);
```

This statement causes two FIFO queues to be created, one for each destination of xStrm. (the queue at the destination of the feedback path will already have been inserted as described in the preceding section). Assuming that xStrm is of type int, then the size of each queue is at least `2*sizeof(int)`, and at the head of each queue at system initialization is the int value 0. This is illustrated graphically in the flow diagram 402 in FIG. 9A. With xStrm thus initialized, the values appearing on the output of the assignment `xStrm += yStrm` are:

```
0
yStrm(0)
yStrm(0) + yStrm(1)
```

```
yStrm(0) + yStrm(1) + yStrm(2)
```

```
·  
·  
·
```

Changing the offset of `aStrm` relative to `bStrm` in a second graphic representation 412 in FIG. 9B is handled in a similar manner. Here, however, *two values* are inserted into the `aStrm` FIFO queue since it may be desirable to offset `aStrm` by two values relative to `bStrm`. That is accomplished using the following stream initialization statement in which 1 and 2 have been chosen as two values 414 to be inserted into `aStrm`'s queue at system initialization.

```
aStrm.initialize(1, 2);
```

The effects of this initialization are illustrated graphically in the representation 412 in FIG. 9B. With `xStrm` thus initialized, the values appearing on the output of the assignment `xStrm += yStrm` are then:

```
1 + bStrm(0)  
2 + bStrm(1)  
aStrm(0) + bStrm(2)  
aStrm(1) + bStrm(3)  
·  
·  
·
```

[00101] As is the case for C variables, all streams must be declared before being used, although certain stream declarations are made implicitly by context, for example, by appearing as module input or output parameters. For explicit stream declarations, the syntax follows that for C variable declarations, but with the declaration now beginning with the keyword `stream`:

```
stream <storage-class-specifier>optional <type> <identifier-list>;
```

Some examples of stream declarations without a storage class specifier are:

```
stream int xStrm, yStrm;
```

```
stream char cStrm;
```

```
stream double dStrm;
```

Of the five storage class specifiers in C – `auto`, `register`, `static`, `extern` and `typedef` – only `static` is permitted in a stream declaration, as in

```
stream static int xStrm, yStrm;
```

[00102] The semantics of a `static`, as well as a non-`static`, stream declaration is determined by the context in which the declaration appears. There are three such contexts, each with its own scope rule. In each case, the stream-declaration scope rule is identical to that of the variable-declaration counterpart. For a stream declaration, with no storage class specifier and appearing inside a module, the declaration scope extends from the declaration to the end of the module. For a stream declaration, with no storage class specifier and appearing outside all modules (and functions), the declaration scope is global – that is, it is visible to the entire program. For a stream declaration, with the `static` storage class specifier and appearing outside all modules (and functions), the declaration scope extends from the declaration to the end of the source file in which the declaration appears.

[00103] Absent from this list are several declaration forms involving storage class specifiers that pertain to variables but not streams. In C, automatic variables, those variables declared with the `auto` storage class specifier or with no specifier at all, lose their values between function invocations. But since streams do their work only within modules and since modules are not invoked (they are always active), automatic streams are an incongruous concept. The `auto` storage class specifier is therefore not applied to stream declarations.

[00104] A variable declaration with the `static` specifier and appearing inside a function indicates that the declared variable retains its value between function calls (function invocations). In the case of modules, however, there is no notion of a call, and so the `static` specifier has no meaning inside a module. The `static` specifier is therefore not used within module scope.

[00105] For variable declarations, the `extern` storage class specifier helps to distinguish those declarations of global variables that act as declarations and definitions from those that act as merely declarations. In the case of streams, however, a declaration is never a definition because a stream declaration never causes storage to be set aside. Storage is allocated only at stream destinations as described below in the section on Stream FIFOs. The `register` and `typedef` storage class specifiers have no relevance to streams and do not appear in stream declarations.

[00106] Stream expressions are the stream counterpart to ordinary C expressions. Apart from substituting input streams for all variables and an output stream for the result, the two types of expressions are very similar. Expressions combine variables and constants to produce new values while stream expressions combine streams and constants to produce new streams. The structure of C expressions and stream expressions are nearly identical. All C operators are valid operators in stream expressions. The same operator precedence applies in both C expressions and stream expressions. C function calls are permitted in stream expressions, just as they are in C expressions. Instantiations of modules with a single output stream are permitted in stream expressions, and are treated similarly to function calls.

[00107] The differences between C expressions and stream expressions lie primarily in when and how they are evaluated. A C expression is evaluated when the thread of control reaches the statement containing the expression. The evaluation is carried out by first replacing each variable by its current value and then performing the requisite operations according to the rules of operator precedence. The value returned by the final operation is then supplied as the evaluation result.

[00108] Unlike evaluation of C expressions, evaluation of stream expressions in the C Stream programming language is not tied to a thread of control. Instead, stream expressions are evaluated opportunistically. As before, evaluation is carried out by performing the requisite operations according to the rules of operator precedence. Instead of substituting values for variables, a value is consumed (popped) from each FIFO queue belonging to an expression input. A FIFO queue is inserted at all stream destinations that are inputs of stream expressions. The evaluation is opportunistic because it is performed whenever there is at least one value in each input FIFO queue of the expression. The result produced by the evaluation, as before, is the value returned by the final operation of the evaluation. That result, however, is handled differently from the C expression case. For a C expression, the use to which the result is put is determined by the context of the expression. For a stream expression, the result is simply placed into the expression's output stream (which may or may not have a name, depending upon whether the expression is an assignment).

[00109] An example of the stream expression may be shown in the following expression in which `xStrm`, `yStrm` and `zStrm` are all streams of type `int`.

$$xStrm * yStrm + 5 * zStrm$$

The values arriving on the three streams begin as follows:

<u>xStrm</u>	<u>yStrm</u>	<u>zStrm</u>
1	2	3
4	5	6
7	8	9
.	.	.
.	.	.
.	.	.

The first three values placed into the (unnamed) output stream of $xStrm * yStrm + 5 * zStrm$ are then:

Output Stream of $xStrm * yStrm + 5 * zStrm$

$$\begin{aligned}
 17 &= 1 * 2 + 5 * 3 \\
 50 &= 4 * 5 + 5 * 6 \\
 101 &= 7 * 8 + 5 * 9 \\
 &\vdots \\
 &\vdots
 \end{aligned}$$

[00110] Among stream expressions, stream assignments are of special interest. There are two types of such stream assignments, the first type has the form

$\langle\text{stream-identifier}\rangle = \langle\text{stream-expression}\rangle$

[00111] Like their C counterparts, assignments to variables, stream assignments of this type have a side effect. In addition to supplying values to its output stream, a stream assignment makes the output of the right-hand-side (RHS) expression a source of the left-hand-side (LHS) stream, and in the process makes the output stream of the RHS expression the output stream of the assignment. The stream assignment also gives a name to the otherwise nameless output stream of the RHS expression. Although a name is not needed for the output stream of a subexpression of a larger expression, a name is essential when the output stream must be directed to a destination outside any enclosing superexpression.

[00112] The stream assignment statement in the following code fragment is an example. A stream expression, assignment or otherwise, becomes a *stream statement* when it is followed by a semicolon.

```

int F(int, int);

int G(int);

stream int M(int xStrm, int yStrm)

{
  .
}
  
```

```

    .
    out = F(xStrm, G(yStrm));
    .
    .
}
```

The expression, `F(xStrm, G(yStrm))` and the subexpression, `G(yStrm)` each have an output stream as stream expressions. In the case of `G(yStrm)`, the output stream is unnamed since the destination of the stream is clear from the context of the expression: the destination is the second input argument of the function `F` in the superexpression `F(xStrm, G(yStrm))`. In the case of the output stream of `F(xStrm, G(yStrm))`, however, a name is required since the destination is outside the expression. That name is assigned in the assignment expression

```
out = F(xStrm, G(yStrm))
```

With this assignment, the output of `F(xStrm, G(yStrm))` becomes a source of `zStrm`, which has a single destination, the output parameter of Module M.

[00113] The second type of stream assignment is of the form

```
(<comma-separated-list-of-stream-identifiers>) = <module-instantiation>
```

It arises when it is desirable to make the outputs of a multi-output module the sources of multiple named streams. To illustrate, the following multi-output module:

```
stream (int, int) tap(int, int, int);
```

If the first output of `tap` is the source of `int` stream `x`, and the second output of `tap` is the source of `int` stream `y`. That is accomplished with the stream assignment

```
(int x, int y) = tap(arg1, arg2, arg3);
```

The assignment makes the i^{th} output of the module a source of the i^{th} stream and gives names to the otherwise nameless output streams of the module.

[00114] Statements within the body of a module fall into two categories (domains), thread and stream. Stream statements deal with streams but not variables. Thread statements deal with variables, and, in a few cases, with streams as well. Statements in the thread domain are mostly C statements, and like C statements, they are imperative (procedural) in nature defining a step-by-step procedure. Associated with such a procedure is a sequential flow of

control, often called a thread, which governs the order in which statements are executed. Stream statements, in contrast, are declarative. Each such statement makes a declaration about the streams appearing in the statement. There is no notion of a step-by-step procedure, as there is in the thread domain, and the order of stream statements within a module body is therefore immaterial with one exception. Just as variables must be declared before being used, so too streams must be declared before being used.

[00115] Because of the nature of the stream domain, there are no counterparts to those C statements that deal with control flow, specifically, `if-else`, `else-if`, `switch`, `for`, `while`, `do-while`, `break`, `continue`, `goto` and `return`. In fact, the only statement type in the stream domain is the stream counterpart to the C expression statement, and, as in C, the most common expression statement is the assignment statement. A stream expression statement has one of the two forms

```
<stream-expression>;  
  
stream <stream-expression>;
```

while a stream assignment statement has one of the two forms

```
<stream-identifier> = <stream-expression>;  
  
stream <stream-identifier> = <stream-expression>;  
  
(<comma-separated-list-of-stream-identifiers>) = <module-instantiation>
```

[00116] An example application using modules, stream initialization, stream declarations, stream expressions and stream statements is a finite-impulse-response (FIR) filter, a commonly used construct in digital signal processing. A FIR filter transforms a discrete-time input signal into a discrete-time output signal. FIG. 10 is a graphical depiction of a 5-tap FIR filter 500 in which $X(z)$ represents a discrete-time input 502 and $Y(z)$ represents a discrete-time output 504. A series of unit delays 506 labeled z^{-1} each cause the incoming discrete-time signal to be delayed by one clock cycle. A series of multipliers 508 each of which multiply the incoming discrete-time signal by a constant coefficient $h(i)$. Finally, a series of adders 510 labeled Σ each of which sums the two incoming signals. The filter 500 is termed a 5-tap filter because 5 delayed versions of the incoming discrete-time signal are each multiplied by a separate coefficient with the 5 resulting products being summed together.

[00117] A discrete-time signal is represented as a stream of samples. The multipliers 508 and adders 510 are each represented as a stream expression. The unit delay is represented by

stream initialization. By initializing one stream with one or more values, the values are offset (delay) in that stream relative to the values in a second stream. This is the principle underlying operation of the UnitDelay module.

```
stream int UnitDelay(int X)
{
    out = X;
    out.initialize(0);
}
```

In the body of UnitDelay, the stream assignment statement

```
out = X;
```

makes X, the input stream of UnitDelay, the source of out, the default output stream of UnitDelay, while the stream initialization statement:

```
out.initialize(0);
```

inserts the value 0 into out at system initialization. This initial value in out has the effect of offsetting (delaying) all subsequent values in out by one value.

[00118] The following is a Stream C implementation of a 5-tap FIR filter such as the filter 500 in FIG. 10 including 10, 20, 30, 40 and 50 as arbitrarily chosen five filter coefficients.

```
stream int UnitDelay(int X)
{
    out = X;
    out.initialize(0);
}

stream (int xOut, int yOut) tap(int xIn, int yIn, int h)
{
    xOut = UnitDelay(xIn);
    yOut = yIn + h*xOut;
}

stream int FIR5(int X)
{
    (int x2, int y2) = tap(X, 10*X, 20);
    (int x3, int y3) = tap(x2, y2, 30);
    (int x4, int y4) = tap(x3, y3, 40);
    (int , out ) = tap(x4, y4, 50);
```

}

[00119] This implementation exhibits concurrency but does so without any explicit concurrency constructs. The concurrency simply falls out from code that, except for the multiple, named outputs of `tap`, resembles ordinary sequential code. In place of variables, there are now streams.

[00120] Each of the four instantiations of `tap` within the body of `FIR5` is computing its own copy of the formula

$$yIn + h*xOut$$

concurrently with the three other instantiations of `tap`. That is made possible by the opportunistic nature of stream expressions and by the continuing arrival of new input values to each of the instantiations of `tap`. Those new values are supplied by seven internal streams of `FIR5`.

`x` conveys values from the input of `FIR5` to inputs of the first tap

`x2` and `y2` convey values from outputs of the first tap to inputs of the second tap

`x3` and `y3` convey values from outputs of the second tap to inputs of the third tap

`x4` and `y4` convey values from outputs of the third tap to inputs of the fourth tap

The `h` input of each instantiation of `tap` is replaced by a constant. That causes the Stream C compiler to replace all instances of `h` within a `tap` instantiation with the constant. All of the computations performed by the instantiations of `tap` are in service to the transformation of `FIR5` input values into `FIR5` output values. Those final output values are supplied by the default output stream of `FIR5`.

`out` conveys values from an output of the fourth tap to the output of `FIR5`

This implementation is an example of how many digital-signal-processing functions are dealt with in Stream C.

[00121] In the FIR-filter example above, the five coefficients, 10, 20, 30, 40, 50, are known at compile time. However, if the `FIR5` coefficients aren't known at compile time or if the coefficients, although constant for extended periods, do change from time to time, another technique needs to be employed. In such a case, these quasi-constants are not true constants

because they do change, and they are not true streams because their values are not consumed (popped from a FIFO queue) by a stream expression or by a thread.

[00122] A quasi-constant stream is similar to an ordinary stream in several respects. It has a type, one or more sources, one or more destinations and a name. It conveys values of the specified type from the specified sources to the specified destinations. However, a quasi-constant stream differs from an ordinary stream in several ways. Where an ordinary stream would have a FIFO queue, a quasi-constant stream has storage for one value of the specified type (much like the storage associated with a variable). The value residing in that storage is neither popped nor consumed when accessed by a stream expression or thread, but instead remains resident in storage. The stored value is updated when a new value enters the stream through one of the stream sources. When that happens, the new value simply overwrites the old value. Because this updating is typically done asynchronously with system operation, the point at which the update is recognized at the stream destination is, in general, indeterminate. The declaration of a quasi-constant stream must specify an initial value to be stored at each stream storage location at system initialization.

[00123] A quasi-constant stream is declared, either in a standalone declaration or in the input or output parameter list of a module, using the following syntax.

```
const <stream-type> <stream-identifier> = <initial-value>
```

The existing C keyword `const`, which ordinarily applies only to variables, indicates that the stream being declared is a quasi-constant stream. (The use of `const` saves having to introduce a new keyword).

[00124] These ideas are illustrated in the following modification of the FIR5 module. Here, the five coefficients 10, 20, 30, 40 and 50 of the original example have been replaced by the five quasi-constant streams `h0`, `h1`, `h2`, `h3` and `h4`. Since the initial values inserted into these streams at system initialization are the same as the original coefficients, the new FIR5 starts operation with the same coefficients as the original. With the new FIR5, however, those coefficients may be updated if conditions warrant.

```
stream int FIR5(int X, const int h0 = 10,
                const int h1 = 20,
                const int h2 = 30,
                const int h3 = 40,
                const int h4 = 50 )
{
    (int x2, int y2) = tap(X, h0*X, h1);
```

```

(int x3, int y3) = tap(x2, y2, h2);
(int x4, int y4) = tap(x3, y3, h3);
(int     , out    ) = tap(x4, y4, h4);
}

```

[00125] FIG. 11A shows a module 600 highlighting a series of FIFO buffers 602 on an input stream 604. FIGs. 11B and C show two additional alternative implementations of the use of the FIFO buffers 602 and the module 600. FIG. 11B shows the use of a FIFO buffer 602 on a series of output streams 606 only. FIG. 11C shows the use of the FIFO buffers 602 on both the input and output streams 604 and 606. From a programmer's perspective the three diagrams in FIGs. 11A-C are identical. From a performance viewpoint having both buffering on input and output as in FIG. 11C allows the module 600 to be scheduled for execution without regard to space being available on the module 600 that is receiving the stream. This does come at a cost of additional memory and an extra scheduling step. The FIFO buffers 602 may reside in virtual memory space, physical memory space, and register file space depending on the implementation.

[00126] An example of a high level scheduling algorithm for the input stream FIFO as in FIG. 11A is as follows.

```

a. Schedule a Module for Execution when {

    There is data available in the input FIFO of the
    input stream(s)

    AND

    There is space available in input stream FIFO of
    the module connected to the current modules
    output stream

}

```

[00127] An example of a high level scheduling algorithm for the output stream FIFO as in FIG. 10B is as follows

```

b. Schedule a Module for Execution when {

    There is data available in the output stream
    FIFO of the module connected to the current
    modules input stream(s)

```

AND

There is space available in FIFO of the output stream(s)

).

[00128] An example of a high level scheduling algorithm for input and output stream FIFO as in FIG. 10C is as follows:

c. Schedule a Module for Execution when {

There is data available in the input FIFO of the input stream(s)

AND

(There is space available in input stream FIFO of the module connected to the current modules output stream

OR

There is space available in FIFO of the output stream(s))

}

THREADS

[00129] Threads provide capabilities crucial to making Stream C a complete and well-rounded language. Threads may appear either within the body of a C function (i.e., a function whose inputs are individual values and whose output is a single value) or within the body of a module (i.e., a function whose inputs and outputs are streams of values). The two types of threads are identical except that threads in the body of a module may, and usually do, access Stream C streams, and for that reason they usually do not terminate. Also, threads in the body of a C function do not access Stream C streams, and like all (well-behaved) C threads they do terminate.

[00130] The distinguishing characteristic of a Stream C thread is its complete decoupling from concurrency issues. There are no concurrency constructs, no direct interactions with other threads and no spawning of new threads. A Stream C thread is thus oblivious to the

fact that, in general, it is operating in a multi-threaded environment. A programmer working in the thread domain can therefore focus on a strictly sequential problem.

[00131] Function declarations and function definitions in Stream C have the same syntax and semantics as their counterparts in C. For a function call in Stream C, the syntax and semantics depend upon whether the call appears in: (a) the body of a function or (b) a stream expression. A Stream C function call in the body of the same (recursive) function or the body of another function has the same syntax and semantics as a regular C function call. A Stream C function call in a stream expression has the same syntax as a C function call but with streams replacing variables in the function-call arguments. The semantics of such a call are similar, but not identical, to those of a regular function call. The differences relate to how each evaluation (call) of the function is performed. More specifically, they concern: (1) how values are obtained for the parameters (streams) appearing in the function-call arguments, (2) the destination of the function-call output, and (3) how control is handled.

[00132] In C, the parameters appearing in the arguments of a function call are all variables, and the value substituted for each such function input variable is the current value of that variable. In Stream C, the parameters appearing in the arguments of a stream-expression function call are all streams, and the value substituted for each such function input stream is either: (a) the value popped (consumed) from the FIFO queue at that stream destination, in the case of a regular stream, or (b) the current value at that stream destination, in the case of a quasi-constant stream.

[00133] In C, the value returned by a function call is passed to the caller. In Stream C, the value returned by a stream-expression function call is placed into the function-call output stream, which may be either named or unnamed. As a stream expression itself, a stream-expression function call always has an output stream. The destinations of the output value are determined by the destinations of the stream.

[00134] In C, a function is called when the thread of control encounters a call to that function. In Stream C, a stream-expression function call is evaluated (i.e., the function is called), without regard to a thread of control. Instead the function is called opportunistically whenever there is at least one value in the FIFO queue of each regular input stream of the function call. Quasi-constant input streams are always prepared to supply a value, and so they never block a function call or evaluation of a stream expression.

[00135] Apart from these three differences, the semantics of regular C function calls and stream-expression function calls are identical. That means that in both cases, the usual thread-based semantics applies to function execution.

[00136] An example of threads in C Stream may be shown with the following definitions of the function GCD and the module GCD4.

```

int GCD(int a, int b)                                // A recursive
function
{
    if ((a >= b) && ((a % b) == 0))      // Beginning of a
thread
    {
        return (b);
    }
    if (a < b)
    {
        return GCD(b, a);                // A function call
    }
    return GCD(b, (a % b));            // A function call
}

stream int GCD4(int w, int x, int y, int z) // A module
{
    out = GCD(GCD(w, x), GCD(y, z)); // A stream expression
with
    }                                // three function
calls

```

[00137] GCD, a classic example of a recursive function, returns the greatest common divisor of two integers. It has two integer inputs, *a* and *b*, and returns an integer result. GCD4, is a module with four integer-stream inputs, *w*, *x*, *y* and *z*, and an integer-stream output. Within the body of GCD4, is the stream-expression statement

```
out = GCD(GCD(w, x), GCD(y, z));
```

and within this statement is the stream expression

```
GCD(GCD(w, x), GCD(y, z))
```

[00138] Since this expression contains the destinations of streams w , x , y and z , there is a FIFO queue at each of those four destinations. Those queues permit the function calls $GCD(w, x)$ and $GCD(y, z)$ to be evaluated (executed) opportunistically and concurrently as described above. Like these two calls, the third call to GCD is performed opportunistically with input values being obtained from the FIFO queues of its two input streams. Those input streams originate as the output streams of the two other calls to GCD , the FIFO queues on those two streams allowing the third call to GCD to be performed concurrently with the first two. The output stream of this third function call is directed by means of the stream assignment to out , the output stream of $GCD4$. This arrangement of function calls to GCD , which is represented as a data-flow graph in FIG. 12, allows data from the four input streams, w , x , y and z , to be streamed through the three concurrently operating function calls to produce a stream of output values, each of which is the greatest common divisor of $w(i)$, $x(i)$, $y(i)$ and $z(i)$, for some integer $i \geq 0$.

[00139] From the stream point of view, it's immaterial how a module transforms input stream values into output stream values. All that matters are the transformation(s) from inputs to outputs (and any side effects). In the examples presented so far, these transformations have been represented in terms of stream expressions, expressions which may be implemented using application-specific hardware, reconfigurable hardware (such as that in FIGs. 1 and 2), a processor executing sequential code or some other mechanism.

[00140] These transformations may be represented explicitly as sequential code residing in the body of a module. Such code executes on a stored-program sequential processor, and exists in what may be referred to as the thread domain. Although the body of a module will typically contain statements exclusively in either the stream domain or the thread domain, that does not preclude the presence of both types of statements in the same module body. In that case, the two domains operate side by side (i.e., concurrently).

[00141] The syntax and semantics of the thread domain are a superset of C as defined informally by Brian W. Kernighan and Dennis M. Ritchie, "C Programming Language," (1978) and formally by the ISO C standard ISO/IEC 9899. The additions to standard C involve operations that allow a thread to access those streams that are visible to the thread, either module input streams, module output streams, streams internal to the module body or global streams. Those stream-access operations are divided into two categories: blocking and non-blocking. To understand those operations, the mechanisms used to regulate the flow of

values in streams and the mechanisms for managing tasks (a task is equivalent to a module instance) are significant as described above with reference to the node wrapper in FIG. 4.

[00142] Flow control and task management are key services provided by the Stream C run-time-support system. Flow control prevents FIFO queue overflow (i.e., the writing of data to a queue that is already full) and FIFO queue underflow (i.e., the reading of data from a queue that is empty). Task management controls when tasks are placed into execution and, in some cases, when task execution is terminated. There are three key elements of the Stream C flow control and task management systems: consumer counts, producer counts, and task managers.

[00143] An integer consumer count is associated with each FIFO queue of a regular (non quasi-constant) stream. All reads by a particular thread of a particular stream access the same FIFO queue and, therefore, the same consumer count. The consumer count sign bit indicates whether the FIFO queue is empty. A sign bit of 1 (the consumer count is negative) indicates that the queue is empty. A sign bit of 0 (the consumer count is non-negative) indicates that the queue is nonempty.

[00144] An integer producer count is associated with each source of each regular (non quasi-constant) stream. The producer count sign bit indicates whether there is space available in all downstream FIFO queues to receive a value inserted at this stream source. A sign bit of 0 (the producer count is non-negative) indicates that not all downstream queues have space to receive a value this output stream. A sign bit of 1 (the producer count is negative) indicates that all downstream queues have space to receive a value this output stream.

[00145] Each processing core such as the nodes 180 in FIG. 2 has a first-in-first-out ready-to-run queue of tasks that have all the resources, including input data, required to begin execution. Each processing core has a task manager responsible for managing the execution of tasks and providing the necessary coordination signals between tasks. A task manager automatically increments the consumer count when data is pushed onto (written to) a FIFO queue, decrements the consumer count when data is popped (consumed) from a FIFO queue, sends backwards acknowledgements to stream sources to indicate that space has become available at destination FIFO queues (the default is to send a backwards acknowledgement after each value is consumed from each FIFO queue). The task manager also increments the producer count of a module output stream when data is written to that stream, decrements the producer count of a module output stream when a backwards acknowledgement is received for that stream, and places a task in the processing core's ready-to-run task queue when the

task has the input data, and any other required resources, it needs to proceed. The task manager places a task into execution when the task is at the head of the ready-to-run task queue and there is an execution unit available and suspends execution of a task when the task does not have the input data it needs to proceed or when the task times out.

[00146] Blocking stream-access operations allow a thread appearing in a module body to access streams visible to the thread such as module input streams, module output streams, streams internal to the module body and global streams. These are the preferred methods for accessing streams because, unlike their non-blocking brethren, they introduce no non-determinism. The blocking and unblocking of such operations is handled automatically by a processing core's task manager.

[00147] There are three such operations, each patterned after a similar operation in C++. The `>>` operator is used to pop (consume) a single value from a stream FIFO queue and assign that value to a variable. It is used in statements of the form

```
<stream-identifier> >> <variable-identifier>;
```

The statement causes a single value to be popped from the stream on the left and assigned to the variable on the right. If, however, the FIFO queue for the stream is empty as indicated by the sign bit of the stream's consumer count, the statement blocks (stalls) and remains blocked until queue becomes nonempty again, as indicated by the sign bit of the stream's consumer count.

[00148] The `<<` operator is used to place the current value of a variable into a stream. It is used in statements of the form

```
<stream-identifier> << <variable-identifier>;
```

The statement causes the current value of the variable on the right to be placed into the stream on the left. If, however, one or more downstream queues do not have space to receive such a value, as indicated by the sign bit of producer count at the stream source, the statement blocks (stalls) and remains blocked until all downstream queues have space to receive the value again, as indicated by the sign bit of the stream's producer count.

[00149] The `peek` operator is used to obtain the value at the head of a stream FIFO queue without popping (consuming). It is used in expressions of the form

```
<stream-identifier>.peek()
```

This expression returns the current value at the head of the FIFO queue of `<stream-identifier>`, but does not pop (consume) that value from the queue. If, however, the FIFO

queue for the stream is empty as indicated by the sign bit of the stream's consumer count, the statement blocks (stalls) and remains blocked until queue becomes nonempty, again, as indicated by the sign bit of the stream's consumer count.

[00150] Like their blocking cousins, non-blocking stream-access operations allow a thread appearing in a module body to access streams visible to the thread such as module input streams, module output streams, streams internal to the module body and global streams. However, unlike blocking operations, non-blocking operations typically involve race conditions that affect the outcome of the operations and therefore introduce non-determinism. There are two such operations:

[00151] An expression of the form

```
<stream-identifier>.consumerCount()
```

returns the consumer count of <stream-identifier>, where <stream-identifier> is a stream that is read by the thread via the `>>` or `peek` operations. This expression is used primarily to test the consumer-count sign bit of <stream-identifier> in order to bypass a `>>` or `peek` operation when the FIFO queue of <stream-identifier> is empty.

[00152] An expression of the form

```
<stream-identifier>.producerCount()
```

returns the producer count of <stream-identifier>, where <stream-identifier> is a stream that is written to by the thread via the `<<` operation. This expression is used primarily to test the producer-count sign bit of <stream-identifier> in order to bypass a `<<` operation when one or more downstream queues do not have space to receive such a new value.

[00153] While threads within module bodies may take many different forms, many will be variations of the following typical form.

```
stream int moduleA(int strm1, ..., int strmN)
{
    int var1, ..., varN, result; // Declare variables
    while true                      // Loop forever
    {
        strm1 >> var1;
        :
        streams
        :
        strmN >> varN;
        :
        :
        // Read values from input
        // Compute result
    }
}
```

```

        .
        out << result;           // Place result into output
stream
    }
}

```

Here, moduleA is a module with one or more input streams and a single output stream. The input and output stream types are arbitrary chosen to be integers. The first thing the thread within the body of moduleA does is declare a variable for each input stream and a variable for the single output stream. The thread then enters an infinite loop in which each iteration involves: (a) reading in (consuming) a value from each input stream, (b) computing (producing) a result and (c) placing that result into the output stream.

ARRAYS

[00154] As in other languages, arrays play an important role in Stream C, but not just arrays of data elements, but also stream arrays and module arrays. Arrays of actual data values (not pointers to arrays of data values) are conveyed concurrently over multiple streams. Stream arrays are especially valuable when used in conjunction with arrays of modules.

[00155] Stream C inherits its syntax and semantics for data arrays from C. That means that when the name of an array is used as (a function) argument, the value passed to the function is the location or address of the beginning of the array, there is no copying of array elements. The same is true for the stream inputs (arguments) and outputs of a module. To illustrate, the GCD4 module from above may be used.

```

stream int GCD4(int w, int x, int y, int z) // A module
with four
{
    arguments                                // integer
    out = GCD(GCD(w, x), GCD(y, z));
}

```

Instead of supplying GCD4 with four separate integer-stream arguments, it is supplied with a single stream argument in which each value is an array of four integers. GCD4 would then be transformed into:

```

stream int GCD4(int* wxyz)                  // A module
with one

```

```

{
    // array
argument
    out = GCD(GCD(wxyz[0], wxyz[1]), GCD(wxyz[2], wxyz[3]));
}

```

[00156] In accordance with C conventions, the single argument of GCD4 is of type `int*`, that is, a pointer to an integer, and in this case, the first integer in an array of four integers. Those four integers are accessed within the body of GCD4 using the standard C operator `[]`. Supplying C-type data arrays to a module is one way to deal with arrays in the context of streams.

[00157] For some applications, supplying a module with a stream of array pointers is insufficient to fully exploit the concurrency inherent in the application. An array of streams, rather than a stream of arrays therefore permits arrays of actual data values, not pointers to arrays of data values, to be conveyed concurrently over multiple streams. Declaring a stream array is identical to declaring a regular C array, with two differences, the keyword `stream` precedes declaration and the size of the array must be known at compile time. This restriction is necessary since, like modules, all streams within an application are instantiated at compile time.

[00158] Examples of stream-array declarations are:

```

stream int array1D[4];

stream int array2D[4][16];

stream int array3D[4][16][9];

```

The first declaration declares `array1D` to be a one-dimensional array of 4 integer streams. Similarly, `array2D` is declared to be two-dimensional array of 64 integer streams, and `array3D` a three-dimensional array of 576 integer streams. Individual streams of a stream array are accessed in the same way as individual elements of a data array. For example,

```
array3D[3][15][7]
```

denotes one of the 576 streams in `array3D`.

[00159] Once a stream array is declared, the entire array, subarrays of the array or individual streams within the array may be referenced. These three cases are illustrated in the following code fragment.

```

stream int moduleA(int);           // A module
declaration
stream int moduleB(int[4]);        //   "   "
stream int moduleC(int[3][4]);      //   "   "
stream int moduleD(int W[3][4])    // A module definition
{
    .
    .
    stream int X = moduleA(W[2][0]); // A stream statement
    stream int Y = moduleB(W[2]);    //   "   "
    stream int Z = moduleC(W);      //   "   "
    .
    .
}

```

Here, declarations for `moduleA`, `moduleB` and `moduleC`, and a partial definition of `moduleD` are shown. The input types of the four modules are:

<u>Module</u>	<u>Input Type</u>
<code>moduleA(int)</code>	A single integer stream
<code>moduleB(int[4])</code>	A 1-dimensional array of 4 integer streams
<code>moduleC(int[3][4])</code>	A 3X4 array of integer streams
<code>moduleD(int[3][4])</code>	A 3X4 array of integer streams

[00160] The input arguments supplied to the instantiations of `moduleA`, `moduleB` and `moduleC` within the body of `moduleD` are as follows.

<u>Module Instantiation</u>	<u>Argument</u>	<u>Argument Type</u>
<code>moduleA(W[2][0])</code>	<code>W[2][0]</code>	A single integer stream
<code>moduleB(W[2])</code>	<code>W[2]</code>	A 1-dimensional array of 4 integer streams
<code>moduleC(W)</code>	<code>W</code>	A 3X4 array of integer streams

In each case, the module-instantiation argument type matches the module input type, and each module instantiation therefore satisfies the strong-typing requirement of Stream C.

[00161] Accessing individual streams of a stream array within a stream expression is also straightforward, as illustrated in this example of a complex-multiply module.

```
stream int[2] complexMult(int X[2], int Y[2])
```

```

{
    out[0] = X[0]*Y[0] - X[1]*Y[1];
    out[1] = X[0]*Y[1] + X[1]*Y[0];
}

}

```

Because operators within stream expressions are concurrently active, the four multiplies, one addition and one subtraction in the stream expressions $X[0]*Y[0] - X[1]*Y[1]$ and $X[0]*Y[1] + X[1]*Y[0]$ are evaluated concurrently.

[00162] Data parallelism, one of the more popular approaches to parallel processing, is a form of parallelism in which the same task is performed concurrently (in parallel) on different pieces of the same data structure, which is typically an array. In Stream C, data parallelism is supported by module arrays.

[00163] A module array is, as its name implies, an array of modules. It is declared by inserting the array dimensions, in square brackets, between the module name and the list of inputs parameters. The following are two examples of module-array declarations:

```

stream int moduleA[3][4](int, int);

stream (int, int) moduleB[3][4](int, int);

```

In both cases, the array dimensions are 3X4.

[00164] Like the definition of an ordinary (standalone) module, the definition of a module array has a body delimited by curly braces ({ and }). The following are two examples of module-array definitions. The first has a single (default) output stream, while the second has two named output streams.

```

stream int moduleA[3][4](int a, int b)
{
    // Module Body

}

stream (int x, int y) moduleB[3][4](int a, int b)
{
    // Module Body
}.

```

[00165] Once a module array is declared (either in a declaration or a definition), the entire array, subarrays of the array or individual modules within the array may be instantiated within a stream statement in the same manner as data arrays and stream arrays. These three cases are illustrated here for `moduleA[3][4]`.

<u>Instance</u>	<u>Instantiations</u>	
<code>moduleA[2][0]</code>	An individual module instantiation	
<code>moduleA[2]</code> (style 1)	A 1-dimensional array of 4 module instantiations	
<code>moduleA[2][]</code> (style 2)	A 1-dimensional array of 4 module instantiations	
<code>moduleA</code> (style 1)	The entire 3X4 array of module instantiations	
<code>moduleA[][]</code> (style 2)	The entire 3X4 array of module instantiations	

[00166] The key attribute of a module array comes to the fore when the array is instantiated at system initialization. Each element of a module array is instantiated as a separate module instantiation, thereby permitting all array elements to operate concurrently. `ModuleA[3][4]` is an example of this concept. When the module is instantiated, 12 (= 3 X 4) separate instantiations of `moduleA` are created, each operating concurrently with the 11 other instantiations. Furthermore, this multiplication of instantiations applies to each instance of `moduleA[3][4]`. Thus if there are three instances of `moduleA[3][4]`, then 36 (=3X12) separate instantiations of `moduleA` are created.

[00167] The personalization of a module-array instantiation determines what data the instantiation operates upon. The instantiation may be personalized by supplying each module instantiation with its own unique data through the instantiation's input streams. The instantiation may also be personalized by allowing each module instantiation to identify its array indices using the `index` operator, thereby enabling the instantiation to access its own unique part of a global array.

[00168] The first type of personalization is illustrated below, where the stream arrays may be used to supply unique data to each element of a module array. The second type of personalization exploits the fact that the array indices of each array-module instantiation are known at compile time. To access those indices, the programmer uses an operator with the following syntax:

```
int index(int i)
```

where i is an integer expression that evaluates to a constant at compile time. At compile time, `index(i)` is replaced the i^{th} index of the instantiation. A compile-time or run-time error occurs if i is outside array bounds.

[00169] Stream arrays and module arrays find their greatest utility when they are coupled using a special array-coupling feature of Stream C. There are three requirements for a coupling: a) the stream array and module array must have the same dimensions; b) the stream array must be connected (coupled) to a module-array input or output; and c) the stream-array type must match the module input/output type.

[00170] When such a coupling occurs, each individual stream in the stream array is connected (coupled) to the input/output stream of the individual module in the module array with the same indices. Thus, if the stream array $S[D_1][D_2] \dots [D_n]$ is coupled to an input/output of the module array $M[D_1][D_2] \dots [D_n]$, then each individual stream $S[i_1][i_2] \dots [i_n]$ is connected to an input/output of the individual module $M[i_1][i_2] \dots [i_n]$ for $0 \leq i_1 < D_1, 0 \leq i_2 < D_2 \dots 0 \leq i_n < D_n$.

[00171] The following is an example of a stream array coupled to the output of one module array and the input of another module array:

```
stream int moduleA[3][2](); // First coupled module
stream void moduleB[3][2](int); // Second coupled module
stream void parentModule()
{
    stream int cStrm[3][2]; // The coupled stream
    cStrm[][] = moduleA[][](); // moduleA output coupled to
    cStrm
    moduleB[][](cStrm[][]); // cStrm coupled to moduleB
    input
}
```

Here, the output stream of `moduleA[3][2]` is coupled to `cStrm[3][2]`, and `cStrm[3][2]` is coupled to the input stream of `moduleB[3][2]`. The two couplings are legal because:

- `cStrm[3][2]`, `moduleA[3][2]` and `moduleB[3][2]` all have the same dimensions.
- `cStrm[3][2]` is connected to an output of `moduleA[3][2]` and an input of `moduleB[3][2]`.

- The type of `cStrm[3][2]`, the output type of `moduleA[3][2]` and the input type of `moduleB[3][2]` are all `int`.

[00172] The following table lists for each individual stream in `cStrm[3][2]`: (a) the module whose output is the stream source, (b) the individual stream in `cStrm[3][2]` and (c) the module whose input is the stream destination.

<u>Stream Source</u>	<u>Stream</u>	<u>Stream Destination</u>
<code>moduleA[0][0]</code>	<code>cStrm[0][0]</code>	<code>moduleB[0][0]</code>
<code>moduleA[0][1]</code>	<code>cStrm[0][1]</code>	<code>moduleB[0][1]</code>
<code>moduleA[1][0]</code>	<code>cStrm[1][0]</code>	<code>moduleB[1][0]</code>
<code>moduleA[1][1]</code>	<code>cStrm[1][1]</code>	<code>moduleB[1][1]</code>
<code>moduleA[2][0]</code>	<code>cStrm[2][0]</code>	<code>moduleB[2][0]</code>
<code>moduleA[2][1]</code>	<code>cStrm[2][1]</code>	<code>moduleB[2][1]</code>

PINGS

[00173] There are situations when a module needs to notify another module that a particular operation, a side effect, performed by the module has been completed. For example, when a module performs an operation on a data structure in global memory, perhaps as one of many modules performing similar operations on the same data structure, that module typically needs to notify a downstream module that the operation has been completed so that a downstream operation or task may be initiated. In these situations, there is no need to return a value, just a signal that a particular task has been completed. For these situations where a signal, but no value, is needed, Stream C provides the `ping` data type. Pings (values of type `ping`) are featureless and completely indistinguishable from one another.

[00174] Pings are used in conjunction with three operators: (1) the `join` operator to provide synchronization of tasks, (2) the `>>` stream-access operator and (3) the `<<` stream-access operator. The first use involves just streams, while the last two uses involve a stream and a thread.

[00175] The `ping` keyword is used to declare one or more streams of type `ping`. For example, the following statement declares that `pStrm0`, `pStrm1` and `pStrm2` are streams of type `ping`:

```
stream ping pStrm0, pStrm1, pStrm2;
```

The ping keyword is also used in a module prototype/definition to declare that a module input or output stream is of type ping, as in:

```
stream ping moduleName(int, ping); .
```

[00176] The first use of pings involves the join operator, which serves to join a ping stream with one or more other streams to produce a single output stream. This operator is similar to the rendezvous operation found in some other computing models. Expressions containing the join operator take one of two forms:

```
<pingStreamArray>.join()
```

```
<pingStream>.join(<streamExpression>)
```

As with all stream expressions, each evaluation of an expression in one of these forms consumes a single value/ping from each input stream and produces a single value/ping on the expression's (unnamed) output stream. If an input stream is empty (devoid of values), evaluation stalls (blocks) until all input streams have at least one value/ping. There is no need for an explicit join operation for non-ping expressions since the effect of a join operation is already subsumed by the semantics of expression evaluation.

[00177] When an expression of the first type is evaluated, a single ping is consumed from each stream in the array of ping streams, and a single ping is emitted on the expression's output stream.

[00178] An evaluation of an expression in the second form entails the consumption of a single ping from <pingStream> and the evaluation of <streamExpression>. The stream expression <streamExpression> may be of arbitrary type, including ping. The value that results from the evaluation of <streamExpression> is emitted on the output stream of the join operation. If the expression is of type ping, the expression evaluates to a single ping. The ping stream thus acts as a gatekeeper – much like the >> operation described above, allowing an evaluation to proceed only when a ping is present in <pingStream>.

[00179] The two forms of the join operation are represented graphically as shown in FIG. 13A and 13B. In FIG. 13A, the individual streams of pingStrm[], a one-dimensional ping-stream array of size n, are joined to produce a single (unnamed) output ping stream. In FIG. 13B, a single ping stream, pingStrm, is joined with the expression expr to produce a single (unnamed) output stream of the same type as expr.

[00180] One example of the join operation may include a Data Structure X on which two operations, Operation A and Operation B, are performed. These operations meet the following requirements: a) neither Operation A nor Operation B is performed except in response to a go signal; b) when a go signal is received, Operation A and Operation B are performed concurrently; and c) before either Operation A or Operation B can be initiated, both operations performed in response to a preceding go signal must be completed.

[00181] A simple solution to this problem employs two instances of the join operator:

```

stream ping moduleA(ping pStrm)
{
    while (true)
    {
        pStrm >> ping;
        // Perform Operation A on Data Structure X
        out << ping;
    }
}

stream ping moduleB(ping pStrm)
{
    while (true)
    {
        pStrm >> ping;
        // Perform Operation B on Data Structure X
        out << ping;
    }
}

stream ping moduleC(ping goStrm)
{
    stream ping startStrm = goStrm.join(doneStrm);
    stream ping StrmA = moduleA(startStrm);
    stream ping StrmB = moduleB(startStrm);
    stream ping doneStrm = StrmA.join(StrmB);
    doneStrm.initialize(ping);
    out = doneStrm;
}

```

moduleA and moduleB encapsulate Operation A and Operation B, respectively. Each has an input ping stream, which initiates one operation per ping, and an output ping stream,

which confirms completion of one operation per ping. moduleC contains one instance of both moduleA and moduleB, and receives go signals via the goStrm input ping stream.

[00182] The six statements in moduleC play the following roles:

```
stream ping startStrm = goStrm.join(doneStrm);
```

joins goStrm and doneStrm to produce startStrm. A ping is thus placed into startStrm when there is a ping on goStrm (i.e., a go signal) and a ping on doneStrm, which, indicates completion of A and B operations performed in response to the preceding go signal.

```
stream ping StrmA = moduleA(startStrm);
```

connects startStrm to the input ping stream of moduleA, and connects the output ping stream of moduleA to StrmA. That means that Operation A is performed in response to a go signal, but only after both operations associated with the preceding go signal have been completed.

```
stream ping StrmB = moduleB(startStrm);
```

is similar to the preceding statement. It ensures that Operation B is performed in response to a go signal, but only after both operations associated with the preceding go signal have been completed. There are, however, no restrictions on the order in which Operations A and B are performed. In other words, they are performed concurrently.

```
stream ping doneStrm = StrmA.join(StrmB);
```

joins StrmA, the output ping stream of moduleA, and StrmB, the output ping stream of moduleB. A ping is thus placed onto doneStrm when both operations performed in response to the preceding go signal have been completed.

```
doneStrm.initialize(ping);
```

places a single ping into doneStrm at system initialization. This indicates that all previous operations, of which there are none, have been completed. Without this statement, moduleC would deadlock and no operations would ever be performed.

```
out = doneStrm;
```

connects doneStrm to out, the default output stream of moduleC. Each ping on this stream confirms that the Operation A and Operation B performed in response to a go signal

have been completed. The behavior of moduleC may be summed up as a go signal (ping) on the input port of moduleC causing Operation A and Operation B to be performed concurrently on Data Structure X, but only after previous operations have been completed. When both Operation A and Operation B are completed, moduleC sends a ping on its output port as confirmation.

[00183] A statement of the form

```
pingStrm >> ping;
```

where pingStrm is a stream of type ping, serves to synchronize execution of a thread with the pings in pingStrm. When the statement is encountered in a thread, a single ping is read (consumed) from pingStrm. If pingStrm is empty (i.e., there are no pings in pingStrm), then the statement blocks (stalls) until a ping becomes available. The statement thus acts as a gatekeeper, allowing a thread to proceed only when a ping is present in pingStrm. There is no variable involved in this operation, on the right of the >> operator, where a variable would ordinarily be expected, is just the keyword ping.

[00184] A statement of the form

```
pingStrm << ping;
```

where pingStrm is a stream of type ping, allows a thread to signal interested parties that a certain operation, or operations, have been completed. When the statement is encountered in a thread, a single ping is written to (placed in) pingStrm. Unlike the first statement above, this statement never blocks.

[00185] These two forms of stream/thread interaction involving pings are illustrated in the following code fragment:

```
stream ping moduleA(ping pStrm)
{
    // Do initialization before entering loop
    while (true)
    {
        pStrm >> ping;

        // Perform operations with side effects

        out << ping;
    }
}
```

moduleA has a single input port and a single output port, both of type ping. Within moduleA is a thread containing an infinite loop, each iteration of which begins with the statement

```
pStrm >> ping;
```

This statement serves to synchronize the iterations of the loop with the pings in the module input stream pStrm. It blocks when pStrm is empty and consumes a single ping from pStrm when pStrm is non-empty. Following that statement are statements associated with an activity that invariably involves side effects. If there were no side effects, moduleA would be equivalent to a no-op. At the end of each iteration is the statement

```
out << ping;
```

which signals through moduleA's standard output port that another loop iteration has been completed.

[00186] The join operator is useful when working entirely within the stream domain. There may be situations, however, in which it is more convenient to do the join within a thread. Consider, for example, joining the individual streams of

```
stream ping pingStrm[32];
```

within a thread. That can be accomplished by embedding a for loop within a thread:

```
for (int i = 0; i < 32; ++i)
{
    pingStrm[i] >> ping;
}
```

This loop blocks a thread until one ping has been consumed from each of the 32 streams in pingStrm. An output stream corresponding to the output stream of pingStrm[].join() is produced by following the for loop with the statement

```
joinStrm << ping;
```

[00187] To create a module that mimics the behavior of pingStrm[].join(), these two code fragments are embedded in a while (true) loop, and the loop is placed in a module:

```
stream ping joinArray(ping pingStrm[32])
{
    while (true)
    {
```

```

        for (int i = 0; i < 32; ++i)
        {
            pingStrm[i] >> ping;
        }
        out << ping;
    }
}

```

[00188] A module with embedded thread may be used to mimic the behavior of `pingStrm.join(expr)`, where `expr` is an expression. In this case, however, the module needs an input stream not only for `pingStrm`, but also for each input stream of `expr`. So, for example, if `expr` is the expression `X*Y + Z`, where `X`, `Y` and `Z` are integers, then the module that implements `pingStrm.join(expr)` looks like:

```

stream ping joinExpr(ping pingStrm, int X, int Y, int Z)
{
    while (true)
    {
        pingStrm >> ping;
        out << X*Y + Z;
    }
}

```

[00189] A pixel-processing example illustrates the use of pings, stream arrays and module arrays in implementing data parallelism, a form of parallelism in which the same task is performed concurrently (in parallel) on different pieces of the same data structure such as an array. The example consists of a module array and a module.

```

extern int xScaleFactor, yScaleFactor;

stream ping doPixel[64][256](int* baStrm) // Body is in the
{                                         // thread domain
    const int x = xScaleFactor * index(0);
    const int y = yScaleFactor * index(1);
    int* baseAddress;
    while true
    {
        baStrm >> baseAddress;
        .
        .
        .
        // Perform computations on
        // baStrm[x][y] and
neighbors
}

```

```

        out << ping;
    }
}

stream void parentModule(int* baStrm)      // Body is in the
{                                         // stream domain
    stream ping xStrm[64][256];
    stream ping jStrm;
    jStrm.initialize(ping);
    xStrm[][] = doPixel[][](jStrm.join(baStrm));
    jStrm      = xStrm[][] .join();
}

```

[00190] The two-dimensional module array, `doPixel[64][256]`, is sized to match the size of a two-dimensional array of pixels. The base addresses of the pixel arrays on which `doPixel[64][256]` operates are supplied by the input stream `baStrm`. The `x` coordinate of the pixels upon an individual `doPixel` module operates is obtained by multiplying `index(0)`, the `x` index of the individual `doPixel` module (see Section 5.3), by the global constant `xScaleFactor`. The `y` coordinate of the pixels upon an individual `doPixel` module operates is obtained by multiplying `index(1)`, the `y` index of the individual `doPixel` module, by the global constant `yScaleFactor`. The processing of each pixel begins by setting the variable `baStrm` to the current value in `baStrm`. What follows are computations performed on `baStrm[x][y]` and its neighbors. When processing is done, the individual `doPixel` module signals completion by emitting a `ping`.

[00191] The `parentModule` is responsible for broadcasting pixel-array base addresses to the individual modules in `doPixel[64][256]`. This is done via the statement:

```
xStrm[][] = doPixel[][](jStrm.join(baStrm));
```

Here, the expression `jStrm.join(baStrm)` in the input argument list of `doPixel` acts as a gate, allowing a value in `baStrm` to pass only when there is a `ping` in `jStrm`. An initial `ping` inserted into `jStrm` by the statement

```
jStrm.initialize(ping);
```

allows the very first base address to pass unimpeded. After that, `pings` are inserted into `jStrm` by the statement

```
jStrm = xStrm[][] .join();
```

where `xStrm[64][256]` is the array of ping streams produced by the individual modules in `doPixel[64][256]`. A new ping is therefore inserted into `jStrm` only when all modules in `doPixel[64][256]` have signaled completion of their previous computation by emitting a ping. This ensures that all computations on a pixel array are completed before computations on the next array are begun.

[00192] There is a significant advantage to using pings rather than a standard C data type. With a C data type, a first-in-first-out queue (FIFO) is needed for data values at every destination of a C-data-type stream, that is, everywhere that the stream is an input to an expression. But because pings are indistinguishable from one another, all that is needed at each destination of a ping stream is a counter to tell the number of pings queued up. This results in a significant cost savings over a first-in-first-out queue for data values.

[00193] Pragmas are directives to the Stream C compiler/linker/loader. The directive `#pragma InitializeCount(m, p, n)` initializes the consumer/producer count of input/output port `p` of module `m` to `n`. The Pragma must immediately follow the module definition `#pragma FwrdsAckValue(m, s, n)`. This definition specifies `n` as the forwards acknowledgement value for the point-to-point connection starting at output stream `s` of module `m`. The Pragma must immediately follow the module definition

`#pragma BwrdsAckValue(m, s, n)` specifies `n` as the backwards acknowledgement value for the point-to-point connection starting at output stream `s` of module `m`. The Pragma must immediately follow the module definition.

[00194] Some example benefits of the above described concepts are support of threads and multi-threading i.e., the concurrent execution of multiple threads. Also, all forms of parallelism are expressible such as SIMD, MIMD, Instruction-Level, Task-Level, Data-Parallel, Data-Flow, and Systolic. Deterministic behavior is the default. Non-determinism is explicitly added to programs, and only where needed, as it is in sequential programming which makes software testability and reliability more efficient. The concepts described above have no explicit parallelism constructs. Parallelism falls out from code in the stream domain that – syntactically, at least – resembles ordinary sequential code. A programmer working in the thread domain can therefore focus on a strictly sequential problem. The programming model lends itself to model-based design and model-based testing and scales to an arbitrary number of processing cores. The programming model is equally applicable whether the distances separating processing cores are measured in nanometers or thousands

of kilometers. There are no foreground or background tasks, just tasks, and there are no interrupts or message passing, just streams.

[00195] Although the invention has been described with respect to specific embodiments, thereof, these embodiments are merely illustrative, and not restrictive of the invention. For example, any type of processing units, functional circuitry or collection of one or more units and/or resources such as memories, I/O elements, etc., can be included in a node. A node can be a simple register, or more complex, such as a digital signal processing system. Other types of networks or interconnection schemes than those described herein can be employed. It is possible that features or aspects of the present invention can be achieved in systems other than an adaptable system, such as described herein with respect to a preferred embodiment.

I/We claim:

1. A programmable core based computing device comprising:
 - a plurality of processing cores coupled to each other;
 - a memory storing stream-domain code including a stream defining a stream destination module and a stream source module, the stream source module placing data values in the stream, the stream conveying data values from the stream source module to the stream destination module; and
 - a runtime system detecting when the data values are available to the stream destination module and scheduling the stream destination module for execution on one of the plurality of processing cores.
2. The device of claim 1, wherein the stream-domain code includes a stream expression including an input stream and an output stream, the stream expression causing a data value to be consumed from the input stream and a data value to be produced in the output stream of the stream expression.
3. The device of claim 2, wherein the stream expression includes a function call corresponding with a function, the stream expression causing the function to be called with the data value retrieved from the input stream and the result returned by the function call to be placed into the output stream, the output stream being associated with the function call.
4. The device of claim 2, wherein the stream-domain code includes an assignment statement causing the output stream of the stream expression to become a source of a second previously declared stream.
5. The device of claim 2, wherein the stream expression is one of a plurality of stream expressions and wherein each of the plurality of stream expressions is executed on a corresponding separate processing core.
6. The device of claim 1, wherein the stream-domain code includes a module having an input stream and an output stream, the module consuming data values from the input stream and producing data values in the output stream.
7. The device of claim 6, wherein the module is one of a plurality of modules, each module being executed on a separate one of the plurality of processing cores.
8. The device of claim 6 wherein the module contains within its body a secondary module.
9. The device of claim 6, wherein the module includes thread-domain code executed sequentially.

10. The device of claim 6, wherein the output stream is one of a plurality of output streams included in the module.
11. The device of claim 6, wherein the module is one module in a plurality of modules organized in a module array contained in the stream-domain code, the module array having a plurality of indices allowing access to each module.
12. The device of claim 11, wherein the stream-domain code includes an array of streams associated with a plurality of array indices, each stream in the array of streams conveying data values from an array of stream sources to an array of stream destinations and each stream accessible via an array index.
13. The device of claim 12, wherein each of the streams of the stream array are coupled to the input streams of the module array or the output streams of the module array.
14. The device of claim 1, further comprising a multi-core runtime system, the runtime system conveying data values from a stream source to a stream destination.
15. The device of claim 14, wherein the runtime system is implemented in one or more of instructions stored on a non-transitory medium that may be executed by a processing core, hardware, or reconfigurable hardware.
16. The device of claim 14, wherein the multi-core runtime system has a first-in-first-out queue at the stream destination.
17. The device of claim 16, wherein the multi-core runtime system prevents the overflow of the first-in-first-out queue.
18. The device of claim 16, further comprising:
 - a run-time system assigning each processing core zero or more tasks to be performed, each task implementing either a stream expression or an instance of thread-domain code appearing in the computer program; and
 - a task manager managing the execution of the tasks assigned to the processing core.
19. The device of claim 18, wherein the task manager maintains a consumer count for each input stream of each task, a producer count for each output stream of each task, a ready-to-run queue of tasks ready to executed, an input count for each task that determines the number of task input streams that are required to be enabled in order for a task to be ready to run, and an output count for each task that determines the number of task output streams that are required to be enabled in order for a task to be ready to run,
wherein the task manager further:

increments the consumer count of an input stream in response to a forward acknowledgement sent from the task at a stream source, said acknowledgement indicating that additional data values have been deposited in the first-in-first-out queue associated with the input stream;

decrements the producer count of an output stream in response to a backward acknowledgement sent from the task at a stream destination, said acknowledgement indicating that data values have been removed from the first-in-first-out queue associated with the input stream at the stream destination;

monitors the enabling of a producer count and a consumer count;

places a task in the ready-to-run task queue when the number of enabled consumer counts reaches the task input count and the number of enabled producer counts reaches the task output count;

places a task into execution when the task is at the head of the ready-to-run task queue and the processor core is available; and

suspends execution of the task if the task completes execution, the task is blocked by stream-access instruction or the task times out

20. The device of claim 18, wherein the task manager is implemented in one or more of instructions stored on a non-transitory medium that may be executed by a processing core, hardware, or reconfigurable hardware.

21. A method of programming a computer system including a plurality of processing cores, the method comprising:

providing a computer program for execution on the plurality of processing cores, the computer programming including stream-domain code, the stream-domain code including a stream conveying data values from at least one stream source corresponding to one of the processing cores to at least one stream destination corresponding to another one of the processing cores.

22. The method of claim 21, wherein the stream-domain code includes a stream expression including an input stream and an output stream, the stream expression causing a data value to be consumed from the input stream and a data value to be produced in the output stream of the stream expression.

23. The method of claim 22, wherein the stream expression includes a function call corresponding with a function, the stream expression causing the function to be called with

the data value retrieved from the input stream and the result returned by the function call to be placed into the output stream, the output stream being associated with the function call.

24. The method of claim 22, wherein the stream-domain code includes an assignment statement causing the output stream of the stream expression to become a source of a second previously declared stream.

25. The method of claim 22, wherein the stream expression is one of a plurality of stream expressions and wherein each of the plurality of stream expressions is executed on a corresponding separate processing core.

26. The method of claim 21, wherein the stream-domain code includes a module having an input stream and an output stream, the module consuming data values from the input stream and producing data values in the output stream.

27. The method of claim 26, wherein the module is one of a plurality of modules, each module being executed on a separate one of the plurality of processing cores.

28. The method of claim 26 wherein the module contains within its body a secondary module.

29. The method of claim 26, wherein the module includes thread-domain code executed sequentially.

30. The method of claim 26, wherein the output stream is one of a plurality of output streams included in the module.

31. The method of claim 26, wherein the module is one module in a plurality of modules organized in a module array contained in the stream-domain code, the module array having a plurality of indices allowing access to each module.

32. The method of claim 31, wherein the stream-domain code includes an array of streams associated with a plurality of array indices, each stream in the array of streams conveying data values from an array of stream sources to an array of stream destinations and each stream accessible via an array index.

33. The method of claim 32, wherein each of the streams of the stream array are coupled to the input streams of the module array or the output streams of the module array.

34. The method of claim 21, further comprising managing execution of a computer program via a multi-core runtime system, the runtime system conveying data values from a stream source to a stream destination.

35. The method of claim 34, wherein the runtime system is implemented in one or more of instructions stored on a non-transitory medium that may be executed by a processing core, hardware, or reconfigurable hardware.

36. The method of claim 34, wherein the multi-core runtime system has a first-in-first-out queue at the stream destination.

37. The method of claim 36, wherein the multi-core runtime system prevents the overflow of the first-in-first-out queue.

38. The method of claim 36, further comprising:

assigning each processing core zero or more tasks to be performed, each task implementing either a stream expression or an instance of thread-domain code appearing in the computer program; and

managing the execution of the tasks assigned to the processing core via a task manager.

39. The method of claim 38, wherein the task manager maintains a consumer count for each input stream of each task, a producer count for each output stream of each task, a ready-to-run queue of tasks ready to be executed, an input count for each task that determines the number of task input streams that are required to be enabled in order for a task to be ready to run, and an output count for each task that determines the number of task output streams that are required to be enabled in order for a task to be ready to run, wherein the task manager further:

increments the consumer count of an input stream in response to a forward acknowledgement sent from the task at a stream source, said acknowledgement indicating that additional data values have been deposited in the first-in-first-out queue associated with the input stream;

decrements the producer count of an output stream in response to a backward acknowledgement sent from the task at a stream destination, said acknowledgement indicating that data values have been removed from the first-in-first-out queue associated with the input stream at the stream destination;

monitors the enabling of a producer count and a consumer count;

places a task in the ready-to-run task queue when the number of enabled consumer counts reaches the task input count and the number of enabled producer counts reaches the task output count;

places a task into execution when the task is at the head of the ready-to-run task queue and the processor core is available; and suspends execution of the task if the task completes execution, the task is blocked by stream-access instruction or the task times out

40. The method of claim 38, wherein the task manager is implemented in one or more of instructions stored on a non-transitory medium that may be executed by a processing core, hardware, or reconfigurable hardware.

41. The method of claim 36 wherein the first-in-first-out queue is initialized with an initial value specified in the computer program.


42. The method of claim 21, wherein the stream-domain code includes a keyword indicating that an identifier denoting a stream.

43. The method of claim 21, wherein the stream-domain code includes streams that are elevated to first-class status permitting the stream to be bound to be an identifier, an input parameter of a function, an output of a function, a parameter within an expression or an output of an expression.

44. The method of claim 43, wherein the function and expression are executed on different processing cores.

45. The method of claim 21, wherein the stream conveys data values from a plurality of stream sources to at least one stream destination, and wherein the data values placed into the stream by the plurality of stream sources are interleaved before reaching each destination.

46. The method of claim 21, wherein the stream-domain code references a quasi-constant stream, the quasi-constant stream conveying data values from a stream source to a stream destination storing at least one data value of a type associated with the quasi-constant stream, and wherein a task at the stream destination uses and reuses the same storing the at least one data value until a new data value inserted at a stream source is conveyed to the stream destination.

**FIG. 1'**

2/11

FIG. 2






FIG. 3



3/11

**FIG. 4**

4/11

FIG. 5





FIG. 6








FIG. 8C





FIG. 8D



8/11

FIG. 9A





FIG. 9B



9/11

FIG. 10





FIG. 11A



10/11

FIG. 11B




FIG. 11C




FIG. 12





FIG. 13A

FIG. 13B



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 11/22152

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(8) - G06F 15/00 (2011.01); G06F 15/76 (2011.01)
 USPC - 712/36

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC(8) - G06F 15/00 (2011.01); G06F 15/76 (2011.01)
 USPC - 712/36

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 USPC - 712/36-37; 717/136

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 PubWEST (PGPB, USPT, EPAB, JPAB)
 Google Scholar
 Search terms used: parallel, concurrent, processor, core, stream, task, variable, parameter, producer, consumer, count, schedule

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2005/0166033 A1 (Jacob) 28 July 2005 (28.07.2005), para. [0038]-[0046], [0107]-[0117]	1, 21, 45, 46
Y		----- 2-20, 22-44
Y	US 2009/0300615 A1 (Andrade et al.) 03 December 2009 (03.12.2009), para. [0017]-[0018], [0035]-[0056], [0070]-[0075]	2-20, 22-41
Y	US 2007/0157166 A1 (Stevens) 05 July 2007 (05.07.2007), para. [0094]-[0098], [0164]-[0168], and [0218]	19, 39
Y	US 2003/0200538 A1 (Ebeling et al.) 23 October 2003 (23.10.2003), para. [0027]-[0028], [0050], [0068]	42-44
A	US 2007/0294512 A1 (Crutchfield et al.) 20 December 2007 (20.12.2007)	1, 21

Further documents are listed in the continuation of Box C.

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 03 March 2011 (03.03.2011)	Date of mailing of the international search report 18 MAR 2011
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774