

(22) Date de dépôt/Filing Date: 2018/10/19

(41) Mise à la disp. pub./Open to Public Insp.: 2019/04/20

(45) Date de délivrance/Issue Date: 2021/03/30

(30) **Priorité/Priority:** 2017/10/20 (GB1717291.7)

(51) Cl.Int./Int.Cl. *G06F 9/46* (2006.01),
G06F 9/38 (2018.01), *G06F 9/52* (2006.01),
H04L 7/00 (2006.01)


(72) Inventeurs/Inventors:
KNOWLES, SIMON CHRISTIAN, GB;
ALEXANDER, ALAN GRAHAM, GB

(73) Propriétaire/Owner:
GRAPHCORE LIMITED, GB

(74) **Agent:** BERESKIN & PARR LLP/S.E.N.C.R.L., S.R.L.

(54) Titre : SYNCHRONISATION D'UN ARRANGEMENT DE TRAITEMENT DE MOSAIQUE MULTIPIECE

(54) Title: **SYNCRONIZATION IN A MULTI-TILE PROCESSING ARRANGEMENT**



(57) **Abrégé/Abstract:**

A processing system comprising multiple tiles and an interconnect between the tiles. The interconnect is used to communicate between a group of some or all of the tiles according to a bulk synchronous parallel scheme, whereby each tile in the group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all tiles in the group have completed the compute phase. Each tile in the group has a local exit state upon completion of the compute phase. The instruction set comprises a synchronization instruction for execution by each tile upon completion of its compute phase to signal a sync request to logic in the interconnect. In response to receiving the sync request from all the tiles in the group, the logic releases the next exchange phase and also makes available an aggregated state of all the tiles in the group.

Abstract

A processing system comprising multiple tiles and an interconnect between the tiles. The interconnect is used to communicate between a group of some or all of the tiles according to a bulk synchronous parallel scheme, whereby each tile in the group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all tiles in the group have completed the compute phase. Each tile in the group has a local exit state upon completion of the compute phase. The instruction set comprises a synchronization instruction for execution by each tile upon completion of its compute phase to signal a sync request to logic in the interconnect. In response to receiving the sync request from all the tiles in the group, the logic releases the next exchange phase and also makes available an aggregated state of all the tiles in the group.

Synchronization in a Multi-Tile Processing Arrangement

Technical Field

The present disclosure relates to synchronizing the workloads of multiple different tiles in a multi-tile processing arrangement, each tile comprising its own processing unit and memory. Particularly, the disclosure relates to bulk synchronous parallel (BSP) communication schemes whereby each of a group of tiles must complete a compute phase before any of the tiles in the group can proceed to an exchange phase.

Background

A multi-threaded processor is a processor which is capable of executing multiple program threads alongside one another. The processor may comprise some hardware that is common to the multiple different threads (e.g. a common instruction memory, data memory and/or execution unit); but to support the multi-threading, the processor also comprises some dedicated hardware specific to each thread.

The dedicated hardware comprises at least a respective context register file for each of the number of threads that can be executed at once. A “context”, when talking about multi-threaded processors, refers to the program state of a respective one of the threads being executed alongside one another (e.g. program counter value, status and current operand values). The context register file refers to the respective collection of registers for representing this program state of the respective thread. Registers in a register file are distinct from general purpose memory in that register addresses are fixed as bits in instruction words, whereas memory addresses can be computed by executing instructions. The registers of a given context typically comprise a respective program counter for the respective thread, and a respective set of operand registers for temporarily holding the data acted upon and output by the respective thread during the computations performed by that thread. Each context may also have a

respective status register for storing a status of the respective thread (e.g. whether it is paused or running). Thus each of the currently running threads has its own separate program counter, and optionally operand registers and status register(s).

One possible form of multi-threading is parallelism. That is, as well as multiple contexts, multiple execution pipelines are provided: i.e. a separate execution pipeline for each stream of instructions to be executed in parallel. However, this requires a great deal of duplication in terms of hardware.

Instead therefore, another form of multi-threaded processor employs concurrency rather than parallelism, whereby the threads share a common execution pipeline (or at least a common part of a pipeline) and different threads are interleaved through this same, shared execution pipeline. Performance of a multi-threaded processor may still be improved compared to no concurrency or parallelism, thanks to increased opportunities for hiding pipeline latency. Also, this approach does not require as much extra hardware dedicated to each thread as a fully parallel processor with multiple execution pipelines, and so does not incur so much extra silicon.

One form of parallelism can be achieved by means of a processor comprising an arrangement of multiple tiles on the same chip (i.e. same die), each tile comprising its own separate respective processing unit and memory (including program memory and data memory). Thus separate portions of program code can be run in parallel on different ones of the tiles. The tiles are connected together via an on-chip interconnect which enables the code run on the different tiles to communicate between tiles. In some cases the processing unit on each tile may itself run multiple concurrent threads on tile, each tile having its own respective set of contexts and corresponding pipeline as described above in order to support interleaving of multiple threads on the same tile through the same pipeline.

In general, there may exist dependencies between the portions of a program running on different tiles. A technique is therefore required to prevent a piece of code on one tile running ahead of data upon which it is dependent being made available by another piece of code on another tile. There are a number of possible schemes for achieving this, but the scheme of interest herein is known as “bulk synchronous parallel” (BSP). According to BSP, each tile performs a compute phase and an exchange phase in an alternating cycle. During the compute phase each tile performs one or more computation tasks locally on tile, but does not communicate any results of its computations with any others of the tiles. In the exchange phase each tile is allowed to exchange one or more results of the computations from the preceding compute phase to and/or from one or more others of the tiles in the group, but does not yet proceed to the next compute phase. Further, according to the BSP principle, a barrier synchronization is placed at the juncture transitioning from the compute phase into the exchange phase, or transitioning from the exchange phase into the compute phase, or both. That is it say, either: (a) all tiles are required to complete their respective compute phases before any in the group is allowed to proceed to the next exchange phase, or (b) all tiles in the group are required to complete their respective exchange phases before any tile in the group is allowed to proceed to the next compute phase, or (c) both. In some scenarios a tile performing computation may be allowed to communicate with other system resources such as a network card or storage disk, as long as no communication with other tiles in the group is involved.

An example use of multi-threaded and/or multi-tiled processing is found in machine intelligence. As will be familiar to those skilled in the art of machine intelligence, a machine intelligence algorithm is based around performing iterative updates to a “knowledge model”, which can be represented by a graph of multiple interconnected nodes. Each node represents a function of its inputs. Some nodes receive the inputs to the graph and some receive inputs from one or more other nodes, whilst the output of some nodes form the inputs of other nodes, and the output of some nodes provide the output of the graph (and in some cases a given node may even have all of these: inputs to the graph, outputs from the graph and connections to other nodes). Further, the function at each node is parameterized by one or more respective

parameters, e.g. weights. During a learning stage the aim is, based on a set of experiential input data, to find values for the various parameters such that the graph as a whole will generate a desired output for a range of possible inputs. Various algorithms for doing this are known in the art, such as a back propagation algorithm based on stochastic gradient descent. Over multiple iterations based on the input data, the parameters are gradually tuned to decrease their errors, and thus the graph converges toward a solution. In a subsequent stage, the learned model can then be used to make predictions of outputs given a specified set of inputs or to make inferences as to inputs (causes) given a specified set of outputs.

The implementation of each node will involve the processing of data, and the interconnections of the graph correspond to data to be exchanged between the nodes. Typically, at least some of the processing of each node can be carried out independently of some or all others of the nodes in the graph, and therefore large graphs expose great opportunities for concurrency and/or parallelism.

Summary

The following describes components of a processor having an architecture which has been developed to address issues arising in the computations involved in machine intelligence applications. The processor described herein may be used as a work accelerator, that is, it receives a workload from an application running on a host computer, the workload generally being in the form of very large data sets to be processed (such as the large experience data sets used by a machine intelligence algorithm to learn a knowledge model, or the data from which to perform a prediction or inference using a previously-learned knowledge model). An aim of the architecture presented herein is to process these very large amounts of data highly efficiently. The processor architecture has been developed for processing workloads involved in machine intelligence. Nonetheless, it will be apparent that the disclosed architecture may also be suitable for other workloads sharing similar characteristics.

When executing different portions of a program over multiple tiles, it may be required to perform a barrier synchronization to bring multiple tiles to a common point of execution. It may also be desired to determine a state of the program as a whole after all tiles have completed a compute phase, e.g. to determine whether or not an exception should be reported to the host, or to make a branch decision to determine whether to branch to a next part of the program or continue iterating the current part. For instance if each of a group of tiles is performing the computations of a respective subgraph of a machine intelligence graph, it may be desired to determine whether the nodes of the subgraph have all satisfied a certain condition indicating that the graph is converging towards a solution. To make such a determination using existing techniques requires a number of steps programmed using general-purpose instructions.

It is recognized herein that it would be desirable to tailor the instruction set of a processor to large-scale multi-threadable applications such as machine learning. According to the present disclosure, this is achieved by providing a dedicated machine code instruction for validating a result of a group of tiles only after all the tiles in the group have completed the current BSP compute phase, thus providing the ability to synchronize the tiles and at the same time determine an overall outcome of the multiple threads with a reduced latency and lower code density.

According to one aspect disclosed herein, there is provided a processing system comprising an arrangement of tiles and an interconnect for communicating between the tiles, wherein:

- each tile comprises an execution unit for executing machine code instructions, each being an instance of a predefined set of instruction types in an instruction set of the processor, each instruction type in the instruction set being defined by a corresponding opcode and zero or more operand fields for taking zero or more operands;

- the interconnect is operable to conduct communications between a group of some or all of the tiles according to a bulk synchronous parallel scheme, whereby each of the tiles in said group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all the tiles in said group have completed the compute

phase, wherein each tile in the group has a local exit state upon completion of the compute phase;

the instruction set comprises a synchronization instruction for execution by each tile in the group upon completion of its compute phase, wherein execution of the synchronization instruction causes the execution unit to send a synchronization request to hardware logic in the interconnect; and

the logic in the interconnect is configured to aggregate the local exit states into a global exit state, and, in response to completion of the compute phase by all the tiles in the group as indicated by receiving the synchronization request from all the tiles in the group, to store the global exit state in a global exit state register on each of the tiles in the group, thereby making the global exit state accessible by a portion of code running on each of the tiles in the group.

According to another aspect disclosed herein, there is provided a processing system comprising an arrangement of tiles and an interconnect for communicating between the tiles, wherein:

each tile comprises a respective execution unit for executing machine code instructions, each being an instance of a predefined set of instruction types in an instruction set of the processor, each instruction type in the instruction set being defined by a corresponding opcode and zero or more operand fields for taking zero or more operands;

the interconnect comprises synchronization logic in the form of dedicated hardware logic for coordinating between a group of some or all of said tiles;

the instruction set comprises a synchronization instruction, wherein the respective execution unit on each respective tile is configured so as if an instance of the synchronization instruction is executed through the respective execution unit, then in response to the opcode of the synchronization instruction to cause an instance of a synchronization request to be transmitted from the respective tile to the synchronization logic in the interconnect, and to suspend instruction issue on the respective tile pending a synchronization acknowledgement being received back from the synchronization logic;

each tile has a register has a local exit state register for storing a local exit state of the tile upon completion of a respective compute phase;

the synchronization logic is configured to aggregate the local exit states of the tiles in said group into a global exit state; and

the synchronization logic is further configured so as, in response to receiving an instance of the synchronization request from all of the tiles of said group, to return the synchronization acknowledgment back to each of the tiles in the group and thereby allow the instruction issue to resume, and to store the global exit state in a global exit state register on each of the tiles in the group, thereby making the global exit state accessible by a portion of code running on each of the tiles in the group

In embodiments the execution unit on each tile may be configured to pause instruction issue in response to executing the synchronization instruction; and the logic in the interconnect may be configured so as, in response to receiving the synchronization request from all the tiles in the group, to send back a synchronization acknowledgment signal to each of the tiles in the group in order to resume the instruction issue.

In embodiments, each of the local exit states and the global exit states may be a single bit.

In embodiments, the aggregation may consist of a Boolean AND of the local exit states or a Boolean OR of the local exit states.

In alternative embodiments, the aggregated exit state may comprise at least two bits representing a trinary value, indicating whether the local exit states are all true, all false or mixed.

In embodiments, each of said group of tiles may comprise a local exit state register arranged to represent the local exit state of the tile.

In embodiments, each tile in the group may comprise:

multiple context register sets, each context register set each arranged to store a program state of a respective one of multiple threads; and

a scheduler arranged to schedule the execution of a respective one of a plurality of worker threads in each of a plurality of time slots in a repeating sequence of interleaved time slots, with the program state of each of the worker threads being stored in a respective one of said context register sets;

wherein according to said bulk synchronous parallel scheme, the exchange phase is held back until all the worker threads on all the tiles in the group have completed the compute phase;

wherein the local exit state on each tile may be an aggregate of an individual exit state output by each of the worker threads on the tile; and

wherein said portion of code may comprise at least one of the multiple threads on the tile.

In embodiments each tile in the group may comprise hardware logic configured to perform said aggregation of the individual exit states into the local exit state.

In embodiments the instruction set may comprise an exit instruction for inclusion in each of the worker threads, and the execution unit may be configured to output the individual exit state of the respective worker thread and terminate the respective worker thread in response to the opcode of the exit instruction.

In embodiments, each of the individual exit states and the local exit states may be a single bit, and the aggregation of the individual exit states may consist of a Boolean AND of the individual exit states or a Boolean OR of the individual exit states.

In embodiments, the local exit state may comprise at least two bits representing a trinary value, indicating whether the individual exit states are all true, all false, or mixed.

In embodiments, the exchange phase may be arranged to be performed by a supervisor thread separate to the worker threads, and said at least one thread may comprise the supervisor thread.

In embodiments, the pausing of the instruction issue may comprise at least pausing issue of instructions from the supervisor thread pending the synchronization acknowledgment.

In embodiments, the context register sets on each tile may comprise: multiple worker context register sets arranged to represent the program state of respective ones of said plurality of worker threads, and an additional supervisor context register set comprising an additional set of registers arranged to represent a program state of the supervisor thread.

In embodiments:

- the supervisor thread may be arranged to begin by running in each of said time slots;

- the instruction set may further comprise a relinquish instruction and the execution unit is configured so as, in response to the opcode of the relinquish instruction, to relinquish the time slot in which the relinquish instruction is executed to the respective worker thread; and

- the exit instruction may cause the respective time slot in which the exit instruction is executed to be handed back to the supervisor thread, such that the supervisor thread resumes running in the respective slot.

In embodiments, said portion of code may be configured to use the global exit state, once valid, to perform a branch decision which depends on the global exit state.

In embodiments, the processing system may be programmed to perform a machine intelligence algorithm in which each node in a graph has one or more respective input edges and one or more respective output edges with the input edges of at least some of the nodes being the output edges of at least some others of the nodes, each node comprising a respective function relating its output edges to its input edges, with each respective function being parameterized

by one or more respective parameters, and each of the respective parameters having an associated error, such that the graph converges toward a solution as the errors in some or all of the parameters reduce; wherein each of the tiles may model a respective subgraph comprising a subset of the nodes in the graph, and each of the local exit states may be used to indicate whether the errors in the one or more parameters of the nodes in the respective subgraph have satisfied a predetermined condition.

In embodiments, said group may be selected at least in part by an operand of the synchronization instruction.

In embodiments, the operand of the synchronization instruction may select whether to include only tiles on the same chip or tiles on different chips in said group.

In embodiments, the operand of the synchronization instruction may select said group from amongst different hierarchical levels of groupings.

In embodiments, the instruction set may further include an abstain instruction, which causes the tile on which the abstain instruction is executed to opt-out of said group.

According to another aspect disclosed herein, there is provided a method of operating a processing system comprising an arrangement of tiles and an interconnect for communicating between the tiles, wherein each tile comprises an execution unit for executing machine code instructions, each being an instance of a predefined set of instruction types in an instruction set of the processor, each instruction type in the instruction set being defined by a corresponding opcode and zero or more operand fields for taking zero or more operands; the method comprising:

conducting communications between a group of some or all of the tiles, via the interconnect, according to a bulk synchronous parallel scheme, whereby each of the tiles in said group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all the tiles in said group have completed the compute phase, wherein each tile in the group has a local exit state upon completion of the compute phase;

wherein the instruction set comprises a synchronization instruction for execution by each tile in the group upon completion of its compute phase, wherein execution of the synchronization instruction causes the execution unit to send a synchronization request to hardware logic in the interconnect; and

the method comprises, in response to completion of the compute phase by all the tiles in the group as indicated by receiving the synchronization request from all the tiles in the group, triggering the logic in the interconnect to aggregate the local exit states into a global exit state, and to store the global exit state in a global exit state register on each of the tiles in the group, thereby making the global exit state accessible by a portion of code running on each of the tiles in the group.

According to another aspect disclosed herein, there is provided a computer program product comprising code embodied on computer readable storage and configured to execute on the processing system of any embodiment disclosed herein, the code comprising a portion for execution on each tile in the group including an instance of the synchronization instruction in each portion.

Brief Description of the Drawings

To aid understanding of the present disclosure and to show how embodiments may be put into effect, reference is made by way of example to the accompanying drawings in which:

Figure 1 is a schematic block diagram of a multi-threaded processing unit,

Figure 2 is a schematic block diagram of a plurality of thread contexts,
Figure 3 schematically illustrates a scheme of interleaved execution time slots,
Figure 4 schematically illustrates a supervisor thread and plurality of worker threads,
Figure 5 is a schematic diagram of logic for aggregating exit states of multiple threads,
Figure 6 schematically illustrates synchronization amongst worker threads on the same tile,
Figure 7 is a schematic block diagram of a processor chip comprising multiple tiles,
Figure 8 is a schematic illustration of a bulk synchronous parallel (BSP) computing model,
Figure 9 is another schematic illustration of a BSP model,
Figure 10 is a schematic illustration of BSP between multi-threaded processing units,
Figure 11 is a schematic block diagram of an interconnect system,
Figure 12 is a schematic illustration of system of multiple interconnected processor chips,
Figure 13 is a schematic illustration of a multi-tier BSP scheme,
Figure 14 is another schematic illustration of a system of multiple processor chips,
Figure 15 is a schematic illustration of a graph used in a machine intelligence algorithm, and
Figure 16 illustrates example wiring for synchronizing between chips.

Detailed Description of Embodiments

The following describes a processor architecture which includes a dedicated instruction in its instruction set for performing a barrier synchronization and at the same time aggregating exit states of multiple threads across multiple tiles into a single aggregated state in an exit state register, wherein this aggregated exit state register exists in every tile and contains the same result for every tile having been aggregated. First however an example processor in which this may be incorporated is described with reference to Figures 1 to 4.

Figure 1 illustrates an example of a processor module 4 in accordance with embodiments of the present disclosure. For instance the processor module 4 may be one tile of an array of like processor tiles on a same chip, or may be implemented as a stand-alone processor on its own chip. The processor module 4 comprises a multi-threaded processing unit 10 in the form of a

barrel-threaded processing unit, and a local memory 11 (i.e. on the same tile in the case of a multi-tile array, or same chip in the case of a single-processor chip). A barrel-threaded processing unit is a type of multi-threaded processing unit in which the execution time of the pipeline is divided into a repeating sequence of interleaved time slots, each of which can be owned by a given thread. This will be discussed in more detail shortly. The memory 11 comprises an instruction memory 12 and a data memory 22 (which may be implemented in different addressable memory unit or different regions of the same addressable memory unit). The instruction memory 12 stores machine code to be executed by the processing unit 10, whilst the data memory 22 stores both data to be operated on by the executed code and data output by the executed code (e.g. as a result of such operations).

The memory 12 stores a variety of different threads of a program, each thread comprising a respective sequence of instructions for performing a certain task or tasks. Note that an instruction as referred to herein means a machine code instruction, i.e. an instance of one of the fundamental instructions of the processor's instruction set, consisting of a single opcode and zero or more operands.

The program described herein comprises a plurality of worker threads, and a supervisor subprogram which may be structured as one or more supervisor threads. These will be discussed in more detail shortly. In embodiments, each of some or all of the worker threads takes the form of a respective "codelet". A codelet is a particular type of thread, sometimes also referred to as an "atomic" thread. It has all the input information it needs to execute from the beginning of the thread (from the time of being launched), i.e. it does not take any input from any other part of the program or from memory after being launched. Further, no other part of the program will use any outputs (results) of the thread until it has terminated (finishes). Unless it encounters an error, it is guaranteed to finish. N.B. some literature also defines a codelet as being stateless, i.e. if run twice it could not inherit any information from its first run, but that additional definition is not adopted here. Note also that not all of the worker threads

need be codelets (atomic), and in embodiments some or all of the workers may instead be able to communicate with one another.

Within the processing unit 10, multiple different ones of the threads from the instruction memory 12 can be interleaved through a single execution pipeline 13 (though typically only a subset of the total threads stored in the instruction memory can be interleaved at any given point in the overall program). The multi-threaded processing unit 10 comprises: a plurality of context register files 26 each arranged to represent the state (context) of a different respective one of the threads to be executed concurrently; a shared execution pipeline 13 that is common to the concurrently executed threads; and a scheduler 24 for scheduling the concurrent threads for execution through the shared pipeline in an interleaved manner, preferably in a round robin manner. The processing unit 10 is connected to a shared instruction memory 12 common to the plurality of threads, and a shared data memory 22 that is again common to the plurality of threads.

The execution pipeline 13 comprises a fetch stage 14, a decode stage 16, and an execution stage 18 comprising an execution unit which may perform arithmetic and logical operations, address calculations, load and store operations, and other operations, as defined by the instruction set architecture. Each of the context register files 26 comprises a respective set of registers for representing the program state of a respective thread.

An example of the registers making up each of the context register files 26 is illustrated schematically in Figure 2. Each of the context register files 26 comprises a respective one or more control registers 28, comprising at least a program counter (PC) for the respective thread (for keeping track of the instruction address at which the thread is currently executing), and in embodiments also a set of one or more status registers (SR) recording a current status of the respective thread (such as whether it is currently running or paused, e.g. because it has encountered an error). Each of the context register files 26 also comprises a respective set of operand registers (OP) 32, for temporarily holding operands of the instructions executed by the

respective thread, i.e. values operated upon or resulting from operations defined by the opcodes of the respective thread's instructions when executed. It will be appreciated that each of the context register files 26 may optionally comprise a respective one or more other types of register (not shown). Note also that whilst the term "register file" is sometimes used to refer to a group of registers in a common address space, this does not necessarily have to be the case in the present disclosure and each of the hardware contexts 26 (each of the register sets 26 representing each context) may more generally comprise one or multiple such register files.

As will be discussed in more detail later, the disclosed arrangement has one worker context register file CX0... CX(M-1) for each of the number M of threads that can be executed concurrently (M=3 in the example illustrated but this is not limiting), and one additional supervisor context register file CXS. The worker context register files are reserved for storing the contexts of worker threads, and the supervisor context register file is reserved for storing the context of a supervisor thread. Note that in embodiments the supervisor context is special, in that it has a different number of registers than each of the workers. Each of the worker contexts preferably have the same number of status registers and operand registers as one another. In embodiments the supervisor context may have fewer operand registers than each of the workers. Examples of operand registers the worker context may have that the supervisor does not include: floating point registers, accumulate registers, and/or dedicated weight registers (for holding weights of a neural network). In embodiments the supervisor may also have a different number of status registers. Further, in embodiments the instruction set architecture of the processor module 4 may be configured such that the worker threads and supervisor thread(s) execute some different types of instruction but also share some instruction types.

The fetch stage 14 is connected so as to fetch instructions to be executed from the instruction memory 12, under control of the scheduler 24. The scheduler 24 is configured to control the fetch stage 14 to fetch an instruction from each of a set of concurrently executing threads in turn in a repeating sequence of time slots, thus dividing the resources of the pipeline 13 into a

plurality of temporally interleaved time slots, as will be discussed in more detail shortly. For example the scheduling scheme could be round-robin or weighted round-robin. Another term for a processor operating in such a manner is a barrel threaded processor.

In some embodiments, the scheduler 24 may have access to one of the status registers SR of each thread indicating whether the thread is paused, so that the scheduler 24 in fact controls the fetch stage 14 to fetch the instructions of only those of the threads that are currently active. In embodiments, preferably each time slot (and corresponding context register file) is always owned by one thread or another, i.e. each slot is always occupied by some thread, and each slot is always included in the sequence of the scheduler 24; though the thread occupying any given slot may happen to be paused at the time, in which case when the sequence comes around to that slot, the instruction fetch for the respective thread is passed over. Alternatively it is not excluded for example that in alternative, less preferred implementations, some slots can be temporarily vacant and excluded from the scheduled sequence. Where reference is made to the number of time slots the execution unit is operable to interleave, or such like, this refers to the maximum number of slots the execution unit is capable of executing concurrently, i.e. the number of concurrent slots the execution unit's hardware supports.

The fetch stage 14 has access to the program counter (PC) of each of the contexts. For each respective thread, the fetch stage 14 fetches the next instruction of that thread from the next address in the program memory 12 as indicated by the program counter. The program counter increments each execution cycle unless branched by a branch instruction. The fetch stage 14 then passes the fetched instruction to the decode stage 16 to be decoded, and the decode stage 16 then passes an indication of the decoded instruction to the execution unit 18 along with the decoded addresses of any operand registers 32 specified in the instruction, in order for the instruction to be executed. The execution unit 18 has access to the operand registers 32 and the control registers 28, which it may use in executing the instruction based on the decoded register addresses, such as in the case of an arithmetic instruction (e.g. by adding, multiplying, subtracting or dividing the values in two operand registers and outputting the

result to another operand register of the respective thread). Or if the instruction defines a memory access (load or store), the load/store logic of the execution unit 18 loads a value from the data memory into an operand register of the respective thread, or stores a value from an operand register of the respective thread into the data memory 22, in accordance with the instruction. Or if the instruction defines a branch or a status change, the execution unit changes value in the program counter PC or one of the status registers SR accordingly. Note that while one thread's instruction is being executed by the execution unit 18, an instruction from the thread in the next time slot in the interleaved sequence can be being decoded by the decode stage 16; and/or while one instruction is being decoded by the decode stage 16, the instruction from the thread in the next time slot after that can be being fetched by the fetch stage 14 (though in general the scope of the disclosure is not limited to one instruction per time slot, e.g. in alternative scenarios a batch of two or more instructions could be issued from a given thread per time slot). Thus the interleaving advantageously hides latency in the pipeline 13, in accordance with known barrel threaded processing techniques.

An example of the interleaving scheme implemented by the scheduler 24 is illustrated in Figure 3. Here the concurrent threads are interleaved according to a round-robin scheme whereby, within each round of the scheme, the round is divided into a sequence of time slots S0, S1, S2..., each for executing a respective thread. Typically each slot is one processor cycle long and the different slots are evenly sized, though not necessarily so in all possible embodiments, e.g. a weighted round-robin scheme is also possible whereby some threads get more cycles than others per execution round. In general the barrel-threading may employ either an even round-robin or a weighted round-robin schedule, where in the latter case the weighting may be fixed or adaptive.

Whatever the sequence per execution round, this pattern then repeats, each round comprising a respective instance of each of the time slots. Note therefore that a time slot as referred to herein means the repeating allocated place in the sequence, not a particular instance of the time slot in a given repetition of the sequence. Put another way, the scheduler 24 apportions

the execution cycles of the pipeline 13 into a plurality of temporally interleaved (time-division multiplexed) execution channels, with each comprising a recurrence of a respective time slot in a repeating sequence of time slots. In the illustrated embodiment, there are four time slots, but this is just for illustrative purposes and other numbers are possible. E.g. in one preferred embodiment there are in fact six time slots.

Whatever the number of time slots the round-robin scheme is divided into, then according to present disclosure, the processing unit 10 comprises one more context register file 26 than there are time slots, i.e. it supports one more context than the number of interleaved timeslots it is capable of barrel-threading.

This is illustrated by way of example in Figure 2: if there are four time slots S0...S3 as shown in Figure 3, then there are five context register files, labelled here CX0, CX1, CX2, CX3 and CXS. That is, even though there are only four execution time slots S0...S3 in the barrel-threaded scheme and so only four threads can be executed concurrently, it is disclosed herein to add a fifth context register file CXS, comprising a fifth program counter (PC), a fifth set of operand registers 32, and in embodiments also a fifth set of one or more status registers (SR). Though note that as mentioned, in embodiments the supervisor context may differ from the others CX0...3, and the supervisor thread may support a different set of instructions for operating the execution pipeline 13.

Each of the first four contexts CX0...CX3 is used to represent the state of a respective one of a plurality of “worker threads” currently assigned to one of the four execution time slots S0...S3, for performing whatever application-specific computation tasks are desired by the programmer (note again this may only be subset of the total number of worker threads of the program as stored in the instruction memory 12). The fifth context CXS however, is reserved for a special function, to represent the state of a “supervisor thread” (SV) whose role it is to coordinate the execution of the worker threads, at least in the sense of assigning which of the worker threads W is to be executed in which of the time slots S0, S1, S2... at what point in the overall program.

Optionally the supervisor thread may have other “overseer” or coordinating responsibilities. For example, the supervisor thread may be responsible for performing barrier synchronisations to ensure a certain order of execution. E.g. in a case where one or more second threads are dependent on data to be output by one or more first threads run on the same processor module 4, the supervisor may perform a barrier synchronization to ensure that none of the second threads begins until the first threads have finished. And/or, the supervisor may perform a barrier synchronization to ensure that one or more threads on the processor module 4 do not begin until a certain external source of data, such as another tile or processor chip, has completed the processing required to make that data available. The supervisor thread may also be used to perform other functionality relating to the multiple worker threads. For example, the supervisor thread may be responsible for communicating data externally to the processor module 4 (to receive external data to be acted on by one or more of the threads, and/or to transmit data output by one or more of the worker threads). In general the supervisor thread may be used to provide any kind of overseeing or coordinating function desired by the programmer. For instance as another example, the supervisor may oversee transfer between the tile local memory 12 and one or more resources in the wider system (external to the array 6) such as a storage disk or network card.

Note of course that four time slots is just an example, and generally in other embodiments there may be other numbers, such that if there are a maximum of M time slots $0 \dots M-1$ per round, the processor module 4 comprises $M+1$ contexts $CX \dots CX(M-1)$ & CXS , i.e. one for each worker thread that can be interleaved at any given time and an extra context for the supervisor. E.g. in one exemplary implementation there are six timeslots and seven contexts.

Referring to Figure 4, the supervisor thread SV does not have its own time slot per se in the scheme of interleaved time slots. Nor do the workers as allocation of slots to worker threads is flexibly defined. Rather, each time slot has its own dedicated context register file ($CX0 \dots CXM-1$) for storing worker context, which is used by the worker when the slot is allocated to the worker, but not used when the slot is allocated to the supervisor. When a given slot is allocated

to the supervisor, that slot instead uses the context register file CXS of the supervisor. Note that the supervisor always has access to its own context and no workers are able to occupy the supervisor context register file CXS.

The supervisor thread SV has the ability to run in any and all of the time slots S0....S3 (or more generally S0...SM-1). The scheduler 24 is configured so as, when the program as a whole starts, to begin by allocating the supervisor thread to all of the time slots, i.e. so the supervisor SV starts out running in all of S0...S3. However, the supervisor thread is provided with a mechanism for, at some subsequent point (either straight away or after performing one or more supervisor tasks), temporarily relinquishing each of the slots in which it is running to a respective one of the worker threads, e.g. initially workers W0...W3 in the example shown in Figure 4. This is achieved by the supervisor thread executing a relinquish instruction, called "RUN" by way of example herein. In embodiments this instruction takes two operands: an address of a worker thread in the instruction memory 12 and an address of some data for that worker thread in the data memory 22:

RUN task_addr, data_addr

The worker threads are portions of code that can be run concurrently with one another, each representing one or more respective computation tasks to be performed. The data address may specify some data to be acted upon by the worker thread. Alternatively, the relinquish instruction may take only a single operand specifying the address of the worker thread, and the data address could be included in the code of the worker thread; or in another example the single operand could point to a data structure specifying the addresses of the worker thread and data. As mentioned, in embodiments at least some of the workers may take the form of codelets, i.e. atomic units of concurrently executable code. Alternatively or additionally, some of the workers need not be codelets and may instead be able to communicate with one another.

The relinquish instruction (“RUN”) acts on the scheduler 24 so as to relinquish the current time slot, in which this instruction is itself executed, to the worker thread specified by the operand. Note that it is implicit in the relinquish instruction that it is the time slot in which this instruction is executed that is being relinquished (implicit in the context of machine code instructions means it doesn’t need an operand to specify this – it is understood implicitly from the opcode itself). Thus the time slot which is given away is the time slot in which the supervisor executes the relinquish instruction. Or put another way, the supervisor is executing in the same space that that it gives away. The supervisor says “run this piece of code at this location”, and then from that point onwards the recurring slot is owned (temporarily) by the relevant worker thread.

The supervisor thread SV performs a similar operation in each of one or more others of the time slots, to give away some or all of its time slots to different respective ones of the worker threads W0...W3 (selected from a larger set W0...wj in the instruction memory 12). Once it has done so for the last slot, the supervisor is suspended (then later will resume where it left off when one of the slots is handed back by a worker W).

The supervisor thread SV is thus able to allocate different worker threads, each performing one or more tasks, to different ones of the interleaved execution time slots S0...S3. When the supervisor thread determines it is time to run a worker thread, it uses the relinquish instruction (“RUN”) to allocate this worker to the time slot in which the RUN instruction was executed.

In some embodiments, the instruction set also comprises a variant of the run instruction, RUNALL (“run all”). This instruction is used to launch a set of more than one worker together, all executing the same code. In embodiments this launches a worker in every one of the processing unit’s slots S0...S3 (or more generally S0.... S(M-1)).

Further, in some embodiments the RUN and/or RUNALL instruction, when executed, also automatically copies some status from one or more of the supervisor status registers CXS(SR) to

a corresponding one or more status registers of the worker thread(s) launched by the RUN or RUNALL. For instance the copied status may comprise one or more modes, such as a floating point rounding mode (e.g. round to nearest or round to zero) and/or an overflow mode (e.g. saturate or use a separate value representing infinity). The copied status or mode then controls the worker in question to operate in accordance with the copied status or mode. In embodiments, the worker can later overwrite this in its own status register (but cannot change the supervisor's status). In further alternative or additional embodiments, the workers can choose to read some status from one or more status registers of the supervisor (and again may change their own status later). E.g. again this could be to adopt a mode from the supervisor status register, such as a floating point mode or a rounding mode. In embodiments however, the supervisor cannot read any of the context registers CX0... of the workers.

Once launched, each of the currently allocated worker threads W0...W3 proceeds to perform the one or more computation tasks defined in the code specified by the respective relinquish instruction. At the end of this, the respective worker thread then hands the time slot in which it is running back to the supervisor thread. This is achieved by executing an exit instruction ("EXIT").

The EXIT instruction takes at least one operand and preferably only a single operand, `exit_state` (e.g. a binary value), to be used for any purpose desired by the programmer to indicate a state of the respective codelet upon ending (e.g. to indicate whether a certain condition was met):

EXIT `exit_state`

The EXIT instruction acts on the scheduler 24 so that the time slot in which it is executed is returned back to the supervisor thread. The supervisor thread can then perform one or more subsequent supervisor tasks (e.g. barrier synchronization and/or exchange of data with external resources such as other tiles), and/or continue to execute another relinquish instruction to allocate a new worker thread (W4, etc.) to the slot in question. Note again therefore that the

total number of threads in the instruction memory 12 may be greater than the number that barrel-threaded processing unit 10 can interleave at any one time. It is the role of the supervisor thread SV to schedule which of the worker threads W0...Wj from the instruction memory 12, at which stage in the overall program, are to be assigned to which of the interleaved time slots S0...SM in the round robin schedule of the scheduler 24.

Furthermore, the EXIT instruction is given a further special function, namely to cause the exit state specified in the operand of the EXIT instruction to be automatically aggregated (by dedicated hardware logic) with the exit states of a plurality of other worker threads being run through the same pipeline 13 of the same processor module 4 (e.g. same tile). Thus an extra, implicit facility is included in the instruction for terminating a worker thread.

An example circuit for achieving this is shown in Figure 5. In this example, the exit states of the individual threads and the aggregated exit state each take the form of a single bit, i.e. 0 or 1. The processor module 4 comprises a register 38 for storing the aggregated exit state of that processor module 4. This register may be referred to herein as the “local consensus” register \$LC (as opposed to a global consensus when the processor module 4 is included as one of an array of similar processor tiles, to be discussed in more detail shortly). In embodiments this local consensus register \$LC 38 is one of the supervisor’s status registers in the supervisor’s context register file CXS. The logic for performing the aggregation comprises an AND gate 37 arranged to perform a logical AND of (A) the exit state specified in the EXIT instructions’ operand and (B) the current value in the local consensus register (\$LC) 38, and to output the result (Q) back into the local consensus register \$LC 38 as a new value of the local aggregate.

At a suitable synchronization point in the program, the value stored in the local consensus register (\$LC) 38 is initially reset to a value of 1. I.e. any threads exiting after this point will contribute to the locally aggregated exit state \$LC until next reset. The output (Q) of the AND gate 37 is 1 if both inputs (A, B) are 1, but otherwise the output Q goes to 0 if any of the inputs (A, B) is 0. Every time an EXIT instruction is executed its exit state is aggregated with those that

have gone before (since last reset). Thus by means of the arrangement shown in Figure 5, the logic keeps a running aggregate of the exit states of any worker threads which have terminated by means of an EXIT instruction since the last time the local consensus register (\$LC) 38 was reset. In this example the running aggregate refers to whether or not all threads so far have exited true: any exit state of 0 from any of the worker threads will mean the aggregate in the register 38 becomes latched to 0 until the next reset. In embodiments the supervisor SV can read the running aggregate at any time by getting the current value from the local consensus register (\$LC) 38 (it does not need to wait for an on-tile synchronization to do so).

The reset of the aggregate in the local consensus register (\$LC) 38 may be performed by the supervisor SV performing a PUT to the register address of the local consensus register (\$LC) 38 using one or more general purpose instructions, in this example to put a value of 1 to the register 38. Alternatively it is not excluded that the reset could be performed by an automated mechanism, for example triggered by executing the SYNC instruction described later herein.

The aggregation circuitry 37, in this case the AND gate, is implemented in dedicated hardware circuitry in the execution unit of the execution stage 18, using any suitable combination of electronic components for forming the functionality of a Boolean AND. Dedicated circuitry or hardware means circuitry having a hard-wired function, as opposed to being programmed in software using general purpose code. The updating of the local exit state is triggered by the execution of the special EXIT instruction, this being one of the fundamental machine code instructions in the instruction set of the processor module 4, having the inherent functionality of aggregating the exit states. Also, the local aggregate is stored in a control register 38, meaning a dedicated piece of storage (in embodiments a single bit of storage) whose value can be accessed by the code running on the pipeline, but which is not usable by the load-store unit (LSU) to store any general purpose data. Instead, the function of data held in a control register is fixed, in this case to the function of storing the locally aggregated exit state. Preferably the local consensus register (\$LC) 38 forms one of the control registers on the processor module 4

(e.g. on tile), whose value the supervisor can access by executing a GET instruction and can be set by executing a PUT instruction.

Note that the circuit shown in Figure 5 is just one example. An equivalent circuit would be to replace the AND gate 37 with an OR gate and to invert the interpretation of the exit states 0 and 1 in software, i.e. $0 \rightarrow \text{true}$, $1 \rightarrow \text{false}$ (with the register 38 being reset to 0 rather than 1 at each synchronization point). Equivalently if the AND gate is replaced with an OR gate but the interpretation of the exit states is not inverted, nor the reset value, then the aggregated state in \$LC will record whether any (rather than all) the worker states exited with state 1. In other embodiments, the exit states need not be single bits. E.g. the exit state of each individual worker may be a single bit, but the aggregated exit state \$LC may comprise two bits representing a trinary state: all workers exited with state 1, all workers exited with state 0, or the workers' exit states were mixed. As an example of the logic for implementing this, one of the two bits encoding the trinary value may be a Boolean AND (or OR) of the individual exit states, and the other bit of the trinary value may be a Boolean OR of the individual exit states. The third encoded case, indicating that the worker's exit states were mixed, can then be formed as the XOR of these two bits.

The exit states can be used to represent whatever the programmer wishes, but one particularly envisaged example is to use an exit state of 1 to indicate that the respective worker thread has exited in a "successful" or "true" state, whilst an exit state of 0 indicates the respective worker thread exited in an "unsuccessful" or "false" state (or vice versa if the aggregation circuitry 37 performs an OR instead of an AND and the register \$LC 38 is reset initially to 0). For instance, consider an application where each worker thread performs a computation having an associated condition, such as a condition indicating whether the error(s) in the one or more parameters of a respective node in the graph of a machine intelligence algorithm has/have fallen within an acceptable level according to a predetermined metric. In this case, an individual exit state of one logical level (e.g. 1) may be used to indicate that the condition is satisfied (e.g. the error or errors in the one or more parameters of the node are within an acceptable level

according to some metric); whilst an individual exit state of the opposite logical level (e.g. 0) may be used to indicate that the condition was not satisfied (e.g. the error or errors are not within an acceptable level according to the metric in question). The condition may for example be an error threshold placed on a single parameter or each parameter, or could be a more complex function of a plurality of parameters associated with the respective computation performed by the worker thread.

As another more complex example, the individual exit states of the workers and the aggregated exit state may each comprise two or more bits, which may be used, for example, to represent a degree of confidence in the results of the worker threads. E.g. the exit state of each individual worker thread may represent a probabilistic measure of confidence in a result of the respective worker thread, and the aggregation logic 37 may be replaced with more complex circuitry for performing a probabilistic aggregation of the individual confidence levels in hardware.

Whatever meaning is given by the programmer to the exit states, the supervisor thread SV can then get the aggregated value from the local consensus register (\$LC) 38 to determine the aggregated exit state of all the worker threads that exited since it was last reset, for example at the last synchronization point, e.g. to determine whether or not all the workers exited in a successful or true state. In dependence on this aggregated value, the supervisor thread may then make a decision in accordance with the programmer's design. The programmer can choose to make whatever use of the locally aggregated exit state that he or she wishes. For example, the supervisor thread may consult the local aggregated exit state in on order to determine whether a certain portion of the program made up of a certain subset of worker threads has completed as expected or desired. If not (e.g. at least one of the worker threads exited in an unsuccessful or false state), it may report to a host processor, or may perform another iteration of the part of the program comprising the same worker threads; but if so (e.g. all the worker threads exited in a successful or true state) it may instead branch to another part of the program comprising one or more new workers.

Preferably the supervisor thread should not access the value in the local consensus register (\$LC) 38 until all the worker threads in question have exited, such that the value stored therein represents the correct, up-to-date aggregate state of all the desired threads. Waiting for this may be enforced by a barrier synchronization performed by the supervisor thread to wait all currently-running local worker threads (i.e. those on the same processor module 4, running through the same pipeline 13) to exit. That is, the supervisor thread resets the local consensus register (\$LC) 38, launches a plurality of worker threads, and then initiates a local barrier synchronization (local to the processing module 4, local to one tile) in order to wait for all the outstanding worker threads to exit before the supervisor is allowed to proceed to get the aggregated exit state from the local consensus register (\$LC) 38.

Referring to Figure 6, in embodiments a SYNC (synchronization) instruction is provided in the processor's instruction set. The SYNC instruction has the effect of causing the supervisor thread SV to wait until all currently executing workers W have exited by means of an EXIT instruction. In embodiments the SYNC instruction takes a mode as an operand (in embodiments its only operand), the mode specifying whether the SYNC is to act only locally in relation to only those worker threads running locally on the same processor module 4, e.g. same tile, as the supervisor as part of which the SYNC is executed on (i.e. only threads through the same pipeline 13 of the same barrel-threaded processing unit 10); or whether instead it is to apply across multiple tiles or even across multiple chips.

SYNC mode // mode $\in$ {tile, chip, zone_1, zone_2}

This will be discussed in more detail later but for the purposes of Figure 6 a local SYNC will be assumed ("SYNC tile", i.e. a synchronization within a single tile).

The workers do not need to be identified as operands of the SYNC instruction, as it is implicit that the supervisor SV is then caused to automatically wait until none of the time slots S0, S1, ... of the barrel-threaded processing unit 10 is occupied by a worker. As shown in Figure 6, once

each of a current batch of workers WLn have all been launched by the supervisor, the supervisor then executes a SYNC instruction. If the supervisor SV launches workers W in all the slots $S0...3$ of the barrel-threaded processing unit 10 (all four in the example illustrated, but that is just one example implementation), then the SYNC will be executed by the supervisor once the first of the current batch of worker threads WLn has exited, thus handing back control of at least one slot to the supervisor SV. Otherwise if the workers do not take up all of the slots, the SYNC will simply be executed immediately after the last thread of the current batch WLn has been launched. Either way, the SYNC causes the supervisor SV to wait for all others of the current batch of workers $WLn-1$ to execute an EXIT before the supervisor can proceed. Only after this the supervisor executes a GET instruction to get the content of the local consensus register (\$LC) 38. This waiting by the supervisor thread is imposed in hardware once the SYNC has been executed. I.e. in response to the opcode of the SYNC instruction, the logic in the execution unit (EXU) of the execution stage 18 causes the fetch stage 14 and scheduler 24 to pause from issuing instructions of the supervisor thread until all outstanding worker threads have executed an EXIT instruction. At some point after getting the value of the local consensus register (\$LC) 38 (optionally with some other supervisor code in between), the supervisor executes a PUT instruction to reset the local consensus register (\$LC) 38 (in the illustrated example to 1).

As also illustrated in Figure 6, the SYNC instruction may also be used to place synchronization barriers between different interdependent layers $WL1, WL2, WL3, \dots$ of worker threads, where one or more threads in each successive layer is dependent on data output by one or more worker threads in its preceding layer. The local SYNC executed by the supervisor thread ensures that none of the worker threads in the next layer $WLn+1$ executes until all the worker threads in the immediately preceding layer WLn have exited (by executing an EXIT instruction).

As mentioned, in embodiments the processor module 4 may be implemented as one of an array of interconnected tiles forming a multi-tile processor, wherein each of tile may be configured as described above in relation to Figures 1 to 6.

This is illustrated further in Figure 7 which shows a single chip processor 2, i.e. a single die, comprising an array 6 of multiple processor tiles 4 and an on-chip interconnect 34 connecting between the tiles 4. The chip 2 may be implemented alone on its own single-chip integrated circuit package, or as one of multiple dies packaged in the same IC package. The on-chip interconnect may also be referred to herein as the “exchange fabric” 34 as it enables the tiles 4 to exchange data with one another. Each tile 4 comprises a respective instance of the barrel-threaded processing unit 10 and memory 11, each arranged as described above in relation to Figures 1 to 6. For instance, by way of illustration the chip 2 may comprise of the order of hundreds of tiles 4, or even over a thousand. For completeness, note also that an “array” as referred to herein does not necessarily imply any particular number of dimensions or physical layout of the tiles 4.

In embodiments each chip 2 also comprises one or more external links 8, enabling the chip- 2 to be connected to one or more other, external processors on different chips (e.g. one or more other instances of the same chip 2). These external links 8 may comprise any one or more of: one or more chip-to-host links for connecting the chip 2 to a host processor, and/or one or more chip-to-chip links for connecting together with one or more other instances of the chip 2 on the same IC package or card, or on different cards. In one example arrangement, the chip 2 receives work from a host processor (not shown) which is connected to the chip via one of the chip-to-host links in the form of input data to be processed by the chip 2. Multiple instances of the chip 2 can be connected together into cards by chip-to-chip links. Thus a host may access a computer which is architected as a single chip processor 2 or as multiple single chip processors 2 possibly arranged on multiple interconnected cards, depending on the workload required for the host application.

The interconnect 34 is configured to enable the different processor tiles 4 in the array 6 to communicate with one another on-chip 2. However, as well as there potentially being dependencies between threads on the same tile 4, there may also be dependencies between the portions of the program running on different tiles 4 in the array 6. A technique is therefore

required to prevent a piece of code on one tile 4 running ahead of data upon which it is dependent being made available by another piece of code on another tile 4.

In embodiments, this is achieved by implementing a bulk synchronous parallel (BSP) exchange scheme, as illustrated schematically in Figures 8 and 9.

According to one version of BSP, each tile 4 performs a compute phase 52 and an exchange phase 50 in an alternating cycle, separated from one to the other by a barrier synchronization 30 between tiles. In the case illustrated a barrier synchronization is placed between each compute phase 52 and the following exchange phase 50. During the compute phase 52 each tile 4 performs one or more computation tasks locally on-tile, but does not communicate any results of these computations with any others of the tiles 4. In the exchange phase 50 each tile 4 is allowed to exchange one or more results of the computations from the preceding compute phase to and/or from one or more others of the tiles in the group, but does not perform any new computations until it has received from other tiles 4 any data on which its task(s) has/have dependency. Neither does it send to any other tile any data except that computed in the preceding compute phase. It is not excluded that other operations such as internal control-related operations may be performed in the exchange phase. In embodiments the exchange phase 50 does not include any non-time-deterministic computations, but a small number of time-deterministic computations may optionally be allowed during the exchange phase 50. Note also that a tile 4 performing computation may be allowed during the compute phase 52 to communicate with other external system resources external to the array of tiles 4 being synchronized – e.g. a network card, disk drive, or field programmable gate array (FPGA) – as long as this does not involve communication with other tiles 4 within the group being synchronized. The communication external to the tile group may optionally utilise the BSP mechanism, but alternatively may not utilize BSP and may instead use some other synchronization mechanism of its own.

According to the BSP principle, a barrier synchronization 30 is placed at the juncture transitioning from the compute phases 52 into the exchange phase 50, or the juncture transitioning from the exchange phases 50 into the compute phase 52, or both. That is to say, either: (a) all tiles 4 are required to complete their respective compute phases 52 before any in the group is allowed to proceed to the next exchange phase 50, or (b) all tiles 4 in the group are required to complete their respective exchange phases 50 before any tile in the group is allowed to proceed to the next compute phase 52, or (c) both of these conditions is enforced. In all three variants it is the individual processors which alternate between phases, and the whole assembly which synchronizes. The sequence of exchange and compute phases may then repeat over multiple repetitions. In BSP terminology, each repetition of exchange phase and compute phase is sometimes referred to as a “superstep” (though note that in the literature the terminology is not always used consistently: sometimes each individual exchange phase and compute phase individually is called a superstep, whereas elsewhere, as in the terminology adopted herein, the exchange and compute phases together are referred to as a superstep).

Note also, it is not excluded that multiple different independent groups of tiles 4 on the same chip 2 or different chips could each form a separate respective BSP group operating asynchronously with respect to one another, with the BSP cycle of compute, synchronize and exchange being imposed only within each given group, but each group doing so independently of the other groups. I.e. a multi-tile array 6 might include multiple internally synchronous groups each operating independently and asynchronously to the other such groups (discussed in more detail later). In some embodiments there is a hierarchical grouping of sync and exchange, as will be discussed in more detail later

Figure 9 illustrates the BSP principle as implemented amongst a group 4i, 4ii, 4iii of some or all of the tiles in the array 6, in the case which imposes: (a) a barrier synchronization from compute phase 52 to exchange phase 50 (see above). Note that in this arrangement, some tiles 4 are allowed to begin computing 52 whilst some others are still exchanging.

According to embodiments disclosed herein, this type of BSP may be facilitated by incorporating additional, special, dedicated functionality into a machine code instruction for performing barrier synchronization, i.e. the SYNC instruction.

In embodiments, the SYNC function takes this functionality when qualified by an inter-tile mode as an operand, e.g. the on-chip mode: SYNC chip.

This is illustrated schematically in Figure 10. In the case where each tile 4 comprises a multi-threaded processing unit 10, then each tile's compute phase 52 may in fact comprise tasks performed by multiple worker threads W on the same tile 4 (and a given compute phase 52 on a given tile 4 may comprise one or more layers WL of worker threads, which in the case of multiple layers may be separated by internal barrier synchronizations using the SYNC instruction with the local on-tile mode as an operand, as described previously). Once the supervisor thread SV on a given tile 4 has launched the last worker thread in the current BSP superstep, the supervisor on that tile 4 then executes a SYNC instruction with the inter-tile mode set as the operand: SYNC chip. If the supervisor is to launch (RUN) worker threads in all the slots of its respective processing unit 10, the "SYNC chip" is executed as soon as the first slot that is no longer needed to RUN any more workers in the current BSP superstep is handed back to the supervisor. E.g. this may occur after the first thread to EXIT in the last layer WL, or simply after the first worker thread to EXIT if there is only a single layer. Otherwise if not all the slots are to be used for running workers in the current BSP superstep, the "SYNC chip" can be executed as soon as the last worker that needs to be RUN in the current BSP superstep has been launched. This may occur once all the workers in the last layer have been RUN, or simply once all the worker threads have been RUN if there is only one layer.

The execution unit (EXU) of the execution stage 18 is configured so as, in response to the opcode of the SYNC instruction, when qualified by the on-chip (inter-tile) operand, to cause the supervisor thread in which the "SYNC chip" was executed to be paused until all the tiles 4 in the array 6 have finished running workers. This can be used to implement a barrier to the next BSP

superstep. I.e. after all tiles 4 on the chip 2 have passed the barrier, the cross-tile program as a whole can progress to the next exchange phase 50.

Figure 11 gives a schematic diagram illustrating the logic triggered by a “SYNC chip” according to embodiments disclosed herein.

Once the supervisor has launched (RUN) all of the threads it intends to in the current compute cycle 52, it then executes a SYNC instruction with the on-chip, inter-tile operand: SYNC chip. This triggers the following functionality to be triggered in dedicated synchronization logic 39 on the tile 4, and in a synchronization controller 36 implemented in the hardware interconnect 34. This functionality of both the on-tile sync logic 39 and the synchronization controller 36 in the interconnect 34 is implemented in dedicated hardware circuitry such that, once the SYNC chip is executed, the rest of the functionality proceeds without further instructions being executed to do so.

Firstly, the on-tile sync logic 39 causes the instruction issue for the supervisor on the tile 4 in question to automatically pause (causes the fetch stage 14 and scheduler 24 to suspend issuing instructions of the supervisor). Once all the outstanding worker threads on the local tile 4 have performed an EXIT, then the sync logic 39 automatically sends a synchronization request “sync_req” to the synchronization controller 36 in the interconnect 34. The local tile 4 then continues to wait with the supervisor instruction issue paused. A similar process is also implemented on each of the other tiles 4 in the array 6 (each comprising its own instance of the sync logic 39). Thus at some point, once all the final workers in the current compute phase 52 have EXITed on all the tiles 4 in the array 6, the synchronization controller 36 will have received a respective synchronization request (sync_req) from all the tiles 4 in the array 6. Only then, in response to receiving the sync_req from every tile 4 in the array 6 on the same chip 2, the synchronization controller 36 sends a synchronization acknowledgement signal “sync_ack” back to the sync logic 39 on each of the tiles 4. Up until this point, each of the tiles 4 has had its supervisor instruction issue paused waiting for the synchronization acknowledgment signal

(sync_ack). Upon receiving the sync_ack signal, the sync logic 39 in the tile 4 automatically unpauses the supervisor instruction issue for the respective supervisor thread on that tile 4. The supervisor is then free to proceed with exchanging data with other tiles 4 via the interconnect 34 in a subsequent exchange phase 50.

Preferably the sync_req and sync_ack signals are transmitted and received to and from the synchronization controller, respectively, via one or more dedicated sync wires connecting each tile 4 to the synchronization controller 36 in the interconnect 34.

Furthermore, in accordance with embodiments disclosed herein, an additional functionality is included in the SYNC instruction. That is, at least when executed in an inter-tile mode (e.g. SYNC chip), the SYNC instruction also causes the local exit states \$LC of each of the synchronized tiles 4 to be automatically aggregated in further dedicated hardware 40 in the interconnect 34. In the embodiments shown this logic takes the form of a multi-input AND gate (one input for each tile 4 in the array 6), e.g. formed from a string of two-input AND gates 40i, 40ii, ... as shown by way of example in Figure 11. This inter-tile aggregation logic 40 receives the value in the local exit state register (local consensus register) \$LC 38 from each tile 4 in the array – in embodiments each a single bit - and aggregates them into a single value, e.g. an AND of all the locally aggregated exit states. Thus the logic forms a globally aggregated exit state across all the threads on all tiles 4 in the array 6.

Each of the tiles 4 comprises a respective instance of a global consensus register (\$GC) 42 arranged to receive and store the global exit state from the global aggregation logic 40 in the interconnect 34. In embodiments this is another of the status registers in the supervisor's context register file CXS. In response to the synchronization request (sync_req) being received from all of the tiles 4 in the array 6, the synchronization controller 36 causes the output of the aggregation logic 40 (e.g. the output of the AND) to be stored in the global consensus register (\$GC) 42 on each tile 4 (it will be appreciated that the "switch" shown in Figure 11 is a schematic representation of the functionality and in fact the updating may be implemented by

any suitable digital logic). This register \$GC 42 is accessible by the supervisor thread SV on the respective tile 4 once the supervisor instruction issue is resumed. In embodiments the global consensus register \$GC is implemented as a control register in the control register file such that the supervisor thread can get the value in the global consensus register (\$GC) 42 by means of a GET instruction. Note that the synchronization logic 36 waits until the sync_req is received from all tiles 4 before updating the value in any of the global consensus registers (\$GC) 42, otherwise an incorrect value may be made accessible to a supervisor thread on a tile that has not yet completed its part of the compute phase 52 and is therefore still running.

The globally aggregated exit state \$GC enables the program to determine an overall outcome of parts of the program running on multiple different tiles 4 without having to individually examine the state of each individual worker thread on each individual tile. It can be used for any purpose desired by the programmer. For instance, in the example shown in Figure 11 where the global aggregate is a Boolean AND, this means that any input being 0 results in an aggregate of 0, but if all the inputs are 1 then the aggregate is 1. I.e. if a 1 is used to represent a true or successful outcome, this means that if any of the local exit states of any of the tiles 4 is false or unsuccessful, then the global aggregated state will also be false or represent an unsuccessful outcome. E.g. this could be used to determine whether or not the parts of the code running on all the tiles have all satisfied a predetermined condition. Thus, the program can query a single register (in embodiments a single bit) to ask “did anything go wrong? Yes or no?” or “have all nodes in the graph reached an acceptable level of error? Yes or no?”, rather than having to examine the individual states of the individual worker threads on each individual tile (and again, in embodiments the supervisor is in fact not able to query the state of the workers except through the exit state registers 38, 42). In other words, the EXIT and SYNC instructions each reduce multiple individual exit states into a single combined state.

In one example use case, the supervisor on one or more of the tiles may report to a host processor if the global aggregate indicated a false or unsuccessful outcome. As another example, the program may perform a branch decision depending on the global exit state. For

example, the program examines the global aggregate exit state \$GC and based on this determines whether to continue looping or whether to branch elsewhere. If the global exit state \$GC is still false or unsuccessful, the program continues iterating the same, first part of the program, but once the global exit state \$GC is true or successful, the program branches to a second, different part of the program. The branch decision may be implemented individually in each supervisor thread, or by one of the supervisors taking on the role of master and instructing the other, slave supervisors on the other tiles (the master role being configured in software).

Note that the aggregation logic 40 shown in Figure 11 is just one example. In another equivalent example, the AND may be replaced with an OR, and the interpretation of 0 and 1 may be inverted ($0 \rightarrow \text{true}$, $1 \rightarrow \text{false}$). Equivalently if the AND gate is replaced with an OR gate but the interpretation of the exit states is not inverted, nor the reset value, then the aggregated state in \$GC will record whether any (rather than all) the tiles exited with locally aggregated state 1. In another example, the global exit state \$GC may comprise two bits representing a trinary state: all tiles' locally aggregated exit states \$LC were state 1, all tiles' locally aggregated exit states \$LC were state 0, or the tiles' locally aggregated exit states \$LC were mixed. As another more complex example, the local exit states of the tiles 4 and the globally aggregated exit state may each comprise two or more bits, which may be used, for example, to represent a degree of confidence in the results of the tiles 4. E.g. the locally aggregated exit state \$LC of each individual tile may represent a statistical, probabilistic measure of confidence in a result of the respective tile 4, and the global aggregation logic 40 may be replaced with more complex circuitry for performing a statistical aggregation of the individual confidence levels in hardware.

As mentioned previously, in embodiments multiple instances of the chip 2 can be connected together to form an even larger array of tiles 4 spanning multiple chips 2. This is illustrated in Figure 12. Some or all of the chips 2 may be implemented on the same IC package or some or all of the chips 2 may be implemented on different IC packages. The chips 2 are connected together by an external interconnect 72 (via the external links 8 shown in Fig. 7). As well as

providing a conduit for exchange of data between tiles 4 on different chips, the external exchange peripheral 72 also provides hardware support for performing barrier synchronization between the tiles 4 on different chips 2 and aggregating the local exit states of the tiles 4 on the different chips 2.

In embodiments, the SYNC instruction can take at least one further possible value of its mode operand to specify an external, i.e. inter-chip, synchronization: SYNC zone_n, wherein zone_n represents an external sync zone. The external interconnect 72 comprises similar hardware logic to that described in relation to Figure 11, but on an external, inter-chip scale. When the SYNC instruction is executed with an external sync zone of two or more chips 2 specified in its operand, this causes the logic in the external interconnect 72 to operate in a similar manner to that described in relation to the internal interconnect 34, but across the tiles 4 on the multiple different chips 2 in the specified sync zone.

That is, in response to an external SYNC, the supervisor instruction issue is paused until all tiles 4 on all chips 2 in the external sync zone have completed their compute phase 52 and submitted a sync request. Further, logic in the external interconnect 72 aggregates the local exit states of all these tiles 4, across the multiple chips 2 in the zone in question. Once all tiles 4 in the external sync zone have made the sync request, the external interconnect 72 signals a sync acknowledgment back to the tiles 4 and stores the cross-chip global aggregate exit state into the global consensus registers (\$GC) 42 of all the tiles 4 in question. In response to the sync acknowledgement, the tiles 4 on all the chips 2 in the zone resume instruction issue for the supervisor.

In embodiments the functionality of the interconnect 72 may be implemented in the chips 2, i.e. the logic may be distributed among the chips 2 such that only wired connections between chips are required (Figures 11 and 12 are schematic).

All tiles 4 within the mentioned sync zone are programmed to indicate the same sync zone via the mode operand of their respective SYNC instructions. In embodiments the sync logic in the external interconnect 72 peripheral is configured such that, if this is not the case due to a programming error or other error (such as a memory parity error), then some or all tiles 4 will not receive an acknowledgement, and therefore that the system will come to a halt at the next external barrier, thus allowing a managing external CPU (e.g. the host) to intervene for debug or system recovery. In other embodiments an error is raised in the case where the sync zones don't match. Preferably however the compiler is configured to ensure the tiles in the same zone all indicate the same, correct sync zone at the relevant time.

Figure 13 illustrates an example BSP program flow involving both internal (on-chip) and external (inter-chip) synchronizations. As shown, it is preferable to keep the internal exchanges 50 (of data between tiles 4 on the same chip 2) separate from the external exchanges 50' (of data between tiles 4 on different chips 2). One reason for this is that a global exchange across multiple chips, which is demarcated by the global sync, may be more “expensive” in terms of latency and load-balancing complexity than for only an on-chip synchronization and exchange. Another possible reason is that exchange of data via the internal (on-chip) interconnect 34 can be made time deterministic, whereas in embodiments exchange of data via the external interconnect 72 may be non-time-deterministic. In such scenarios it may be useful to separate internal and external exchange so that the external sync & exchange process does not “contaminate” the internal sync & exchange.

Accordingly, to achieve such separation, in embodiments the program is arranged to perform a sequence of synchronizations, exchange phases and compute phases comprising in the following order: (i) a first compute phase, then (ii) an internal barrier synchronization 30, then (iii) an internal exchange phase 50, then (iv) an external barrier synchronization 80, then (v) an external exchange phase 50'. See chip 2II in Figure 13. The external barrier 80 is imposed after the internal exchange phase 50, such that the program only proceeds to the external exchange 50' after the internal exchange 50. Note also that as shown with respect to chip 2I in Figure 12,

optionally a compute phase may be included between internal exchange (iii) and external barrier (iv). The overall sequence is enforced by the program (e.g. being generated as such by the compiler), and the internal synchronization and exchange does not extend to any tiles or other entities on another chip 2. The sequence (i)-(v) (with the aforementioned optional compute phase between iii and iv) may be repeated in a series of overall iterations. Per iteration there may be multiple instances of the internal compute, sync and exchange (i)-(iii) prior to the external sync & exchange.

Note that during an external exchange 50 the communications are not limited to being only external: some tiles may just perform internal exchanges, some may only perform external exchanges, and some may perform a mix. Note also that as shown in Figure 13, it is in general possible to have a null compute phase 52 or a null exchange phase 50 in any given BSP superstep.

In some embodiments, also as shown in Figure 13, some tiles 4 may perform local input/output during a compute phase, for example they may exchange data with a host.

As illustrated in Figure 14, in embodiments the mode of the SYNC instruction can be used to specify one of multiple different possible external sync zones, e.g. zone_1 or zone_2. In embodiments these correspond to different hierarchical levels. That is to say, each higher hierarchical level 92 (e.g. zone 2) encompasses two or more zones 91A, 91B of at least one lower hierarchical level. In embodiments there are just two hierarchical levels, but higher numbers of nested levels are not excluded. If the operand of the SYNC instruction is set to the lower hierarchical level of external sync zone (SYNC zone_1), then the above-described sync and aggregation operations are performed in relation to the tiles 4 on the chips 2 in only the same lower-level external sync zone as the tile on which the SYNC was executed. If on the other hand the operand of the SYNC instruction is set to the higher hierarchical level of external sync zone (SYNC zone_2), then the above-described sync and aggregation operations are automatically performed in relation to all the tiles on all the chips 2 in the same higher-level

external sync zone as the tile on which the SYNC was executed. In embodiments the highest hierarchical level of sync zone encompasses all all chips, i.e. it is used to perform a global sync. When multiple lower-level zones are used, BSP may be imposed internally amongst the group of tiles 4 on the chip(s) 2 within each zone, but each zone may operate asynchronously with respect to one another until a global sync is performed.

Note that in other embodiments, the sync zones that can be specified by the mode of the SYNC instruction are not limited to being hierarchical in nature. In general, a SYNC instruction may be provided with modes corresponding to any kind of grouping. For instance, the modes may enable selection from amongst only non-hierarchical groups, or a mixture of hierarchical groupings and one or more non-hierarchical groups (where at least one group is not entirely nested within another). This advantageously enables the flexibility for the programmer or compiler, with minimal code density, to select between different layouts of internally-synchronous groups that are asynchronous with respect to one another.

An example mechanism for implementing the synchronization amongst the selected sync group 91, 92 is illustrated in Figure 16. As illustrated, the external sync logic 76 in the external interconnect 72 comprises a respective sync block 95 associated with each respective chip 2. Each sync block 95 comprises respective gating logic and a respective sync aggregator. The gating logic comprises hardware circuitry which connects together the chips 2 in a daisy chain topology for the purpose of synchronization and exit state aggregation, and which propagates the sync and exit state information in accordance with the following. The sync aggregator comprises hardware circuitry configured to aggregate the synchronization requests (sync_req) and the exit states in accordance with the following.

The respective sync block 95 associated with each chip 2 is connected to its respective chip 2, such that it can detect the sync request (Sync_req) raised by that chip 2 and the exit state of that chip 2, and so that it can return the sync acknowledgment (Sync_ack) and global exit state to the respective chip 2. The respective sync block 95 associated with each chip 2 is also

connected to the sync block 95 of at least one other of the chips 2 via an external sync interface comprising a bundle of four sync wires 96, details of which will be discussed in more detailed shortly. This may be part of one of the chip-to-chip links 8. In the case of a link between chips 2 on different cards, the interface 8 may for example comprise a PCI interface and the four sync wires 96 may be implemented by re-using four wires of the PCI interface. Some of the chips' sync blocks 95 are connected to that of two adjacent chips 2, each connection via a respective instance of the four sync wires 96. This way, the chips 2 can be connected in one or more daisy chains via their sync blocks 95. This enables the sync requests, sync acknowledgments, running aggregates of exit states, and global exit states, to be propagated up and down the chain.

In operation, for each sync group 91, 92, the sync block 95 associated with one of the chips 2 in that group is set as the master for synchronization and exit state aggregation purposes, the rest in the group being slaves for this purpose. Each of the slave sync blocks 95 is configured with the direction (e.g. left or right) that it needs to propagate sync requests, sync acknowledgments and exit states for each sync group 91, 92 (i.e. the direction toward the master). In embodiments these settings are configurable by software, e.g. in an initial configuration phase after which the configuration remains set throughout the subsequent operation of the system. For instance this may be configured by the host processor. Alternatively it is not excluded that the configuration could be hard-wired. Either way, the different sync groups 91, 92 can have different masters and in general it is possible for a given chip 2 (or rather its sync block 95) to be master of one group and not another group of which it is a member, or to be master of multiple groups.

For instance, by way of illustration consider the example scenario of Figure 16. Say for the sake of example that the sync block 95 of chip 2IV is set as the master of a given sync group 91A. Consider now the first chip 2I in the chain of chips 2, connected via their sync blocks 95 and wires 96 ultimately to chip 2IV. When all the worker threads of the current compute phase on the first chip 2I have executed an EXIT instruction, and the supervisors on all the (participating) tiles 4 have all executed a SYNC instruction specifying the sync group 91A, then the first chip 2I

signals its sync readiness to its respective associated sync block 95. The chip 2I also outputs to its respective sync block 95 its chip-level aggregated exit state (the aggregate of all the exiting workers on all the participating tiles on the respective chip 2I). In response, the sync block 95 of the first chip 2I propagates a sync request (Sync_req) to the sync block 95 of the next chip 2II in the chain. It also propagates the exit state of the first chip 2I to the sync block 95 of this next chip 2II. The sync block 95 of this second chip 2II waits until the supervisors of its own (participating) tiles 4 have all executed a SYNC instruction specifying the sync group 91A, causing the second chip 2II to signal sync readiness. Only then does the second chip's sync block 95 propagate a sync request to the sync block 95 of the next (third) chip 2III in the chain, and also propagates a running aggregate of the exit state of the first chip 2I with that of the second 2II. If the second chip 2II had become sync ready before the first 2I, then the sync block 95 of the second chip 2II would have waited for the first chip 2I to signal a sync request before propagating the sync request to the sync block 95 of the third chip 2III. The sync block 95 of the third chip 2III behaves in a similar manner, this time aggregating the running aggregate exit state from the second chip 2II to obtain the next running aggregate to pass onwards, etc. This continues toward the master sync block, that of chip 2IV in this example.

The sync block 95 of the master then determines a global aggregate of all the exit states based on the running aggregate it receives and the exit state of its own chip 2IV. It propagates this global aggregate back out along the chain to all the chips 2, along with the sync acknowledgement (Sync_ack).

If the master is part way along a chain, as opposed to being at one end as in the above example, then the sync and exit state information propagates in opposite directions either side of the master, both sides toward the master. In this case the master only issues the sync acknowledgement and global exit state once the sync request from both sides has been received. E.g. consider the case where chip 2III is master of group 92. Further, in embodiments the sync block 95 of some of the chips 2 could connect to that of three or more other chips 2, thus creating multiple branches of chains toward the master. Each chain then behaves as

described above, and the master only issues the sync acknowledgment and global exit state once the sync request from all chains has been received. And/or, one or more of the chips 2 could connect to an external resource such as the host processor, a network card, a storage device or an FPGA.

In embodiments the signalling of the sync and exit state information is implemented as follows. The bundle of four sync wires 96 between each pair of chips 2 comprises two pairs of wires, a first pair 96_0 and a second pair 96_1. Each pair comprises an instance of a sync request wire and an instance of a sync acknowledgment wire. To signal a running aggregate exit state of value 0, the sync block 95 of the sending chip 2 uses the sync request wire of the first wire pair 96_0 when signalling the sync request (sync_req), or to signal a running aggregate of value 1 the sync block 95 uses the sync request wire of the second wire pair 96_1 when signalling the sync request. To signal a global aggregate exit state of value 0, the sync block 95 of the sending chip 2 uses the sync acknowledgment wire of the first wire pair 96_0 when signalling the sync acknowledgment (sync_ack), or to signal a global aggregate of value 1 the sync block 95 uses the sync request wire of the second wire pair 96_1 when signalling the sync acknowledgment.

Note that the above is only the mechanism for propagating sync and exit state information. The actual data (content) is transmitted by another channel, for example as discussed later with reference to Figure 16. Further, it will be appreciated that this is only one example implementation, and the skilled person will be capable of building other circuits for implementing the disclosed synchronization and aggregation functionality once given the specification of that functionality disclosed herein. For instance, the synchronisation logic (95 in Fig 18) could instead use packets carried over the interconnect 34, 72 as an alternative to dedicated wiring. E.g. the sync_req and/or the sync_ack could each be transmitted in the form of one or more packets.

The functionality of the SYNC instruction in the different possible modes is summarized below.

SYNC tile (performs local, on-tile barrier synchronization)

- Supervisor run mode transitions from executing to waiting for workers to exit
- Suspend instruction issue for the supervisor thread until all worker threads are inactive
- When all worker thread are inactive, the aggregated worker exit status is made available via the local consensus register (\$LC) 38.

SYNC chip (performs an internal, on-chip barrier synchronization)

- Supervisor run mode transitions from executing to waiting for workers to exit
- Suspend instruction issue for the supervisor thread until all worker threads are inactive
- When all worker threads are inactive:
 - the aggregated local worker exit status is made available via the local consensus register (\$LC) 38
 - internal sync participation is signalled to the exchange fabric 34
 - supervisor remains inactive until tile 4 receives internal sync acknowledgement from the exchange fabric 34
 - the system-wide exit status is updated in the global consensus register (\$GC) 42.

SYNC zone_n (performs an external barrier synchronization across zone n)

- Supervisor run mode transitions from executing to waiting for workers to exit
- Suspend instruction issue for the Supervisor thread until all Worker threads are Inactive.
- When all worker threads are inactive:
 - the aggregate local worker exit status is available via the local consensus register (\$LC) 38
 - external sync participation is signalled to the external system, e.g. the sync logic in the aforementioned external interconnect 72
 - supervisor remains suspended until tile 4 receives external sync acknowledgement from the external system 72

the system-wide exit status is updated in the global consensus register (\$GC) 42.

As mentioned previously, not all tiles 4 need necessarily participate in the synchronization. In embodiments, as discussed, the group of participating tiles can be set by the mode operand of the sync instruction. However, this only allows for selection of predefined groups of tiles. It is recognized herein that it would also be desirable to be able to select sync participation on a tile-by-tile basis. Therefore in embodiments, there is provided an alternative or additional mechanism for selecting which individual tiles 4 participate in the barrier synchronization.

Particularly, this is achieved by providing an additional type of instruction in the processor instruction set, to be executed by one or some tiles 4 in place of the SYNC instruction. This instruction may be referred to as the “abstain” instruction, or “SANS” instruction (start automatic non-participatory sync). In embodiments the SANS is reserved for use by the supervisor thread. In embodiments it takes a single immediate operand:

SANS n_barriers

The behaviour of the SANS instruction is to cause the tile on which it is executed to abstain from the current barrier synchronization, but without holding up the other tiles which are waiting for all tiles in the specified sync group to SYNC. In effect it says “go on without me”. When the SANS instruction is executed, the opcode of the SANS instruction triggers the logic in the execution unit of the execution stage 18 to send an instance of the synchronization request signal (Sync_req) to the internal and/or external sync controller 36, 76 (depending on the mode). In embodiments, the sync request generated by the SANS applies to any sync group 91, 92 that encompass the tile 4 that executed the SANS. I.e. for whatever sync group the tiles 4 in this local chip or chips are using next (they must agree on the sync group), the sync_req from those that have executed SANS will always be valid.

Thus from the perspective of the sync controller logic 36, 76 and the other tiles 4 in the sync group, the tile 4 executing the SANS instruction appears exactly as a tile 4 executing a SYNC instruction, and does not hold up the sync barrier and the sending of the sync acknowledgment

signal (Sync_ack) from the sync logic 36, 76. I.e. the tiles 4 executing the SANS instead of the SYNC does not hold up or stall any of the other tiles 4 involved any sync group of which the tile in question is otherwise a member. Any handshake performed by a SANS is valid for all sync groups 91, 92.

However, unlike the SYNC instruction, the SANS instruction does not cause supervisor instruction issue to be paused awaiting the sync acknowledgment signal (Sync_ack) from the sync logic 36, 76. Instead the respective tile can simply continue uninhibited by the current barrier synchronization that is being conducted between the other tiles 4 that executed SYNC instructions. Thus by mimicking a sync but not waiting, the SANS instruction allows its tile 4 to press on with processing one or more tasks whilst still allowing the other tiles 4 to sync.

The operand n_barriers specifies the number of “posted” syncs, i.e. the number of future sync points (barriers) the tile will not be participating in. Alternatively it is not excluded that in other embodiments the SANS instruction does not take this operand, and instead each execution of the SANS instruction causes only a one-off abstention.

By means of the SANS instruction, certain tiles 4 may be responsible for performing tasks outside the direct scope of the BSP operating schedule. For example, it may be desirable to allocate a small number of tiles 4 within a chip 2 to initiate (and process) data transfers to and/or from host memory while the majority of tiles 4 are occupied with the primary computation task(s). In such scenarios those tiles 4 not directly involved with primary computation can declare themselves as effectively disconnected from the synchronization mechanism for a period of time using the automatic non-participatory sync feature (SANS). When using this feature, a tile 4 is not required to actively (i.e. via execution of the SYNC instruction) signal its readiness for synchronization (for any of the synchronization zones), and in embodiments makes a null contribution to the aggregated exit status.

The SANS instruction begins or extends a period during which the tile 4 on which it is executed will abstain from active participation in inter-tile synchronization (or synchronization with other external resources if they are also involved in the synchronization). During this period, this tile 4 will automatically signal its readiness for synchronization, within all zones, and in embodiments also make a null contribution to the global aggregated consensus \$GC. This time period may be expressed as an unsigned immediate operand (`n_barriers`) indicating how many additional future synchronization points will be automatically signalled by this tile 4. Upon execution of the SANS, the value `n_barriers` specified by its operand is placed into a countdown register \$ANS_DCOUNT on the respective tile 4. This is a piece of architectural state used to keep a track of how many additional future `sync_reqs` should be made. If the automatic non-participatory sync mechanism is currently inactive, the first assertion of readiness (`sync request, sync_req`) will be performed immediately. Subsequent assertions will occur in the background, once the previous synchronization has completed (i.e. following assertion of the `sync acknowledgment, sync_ack`). If the automatic non-participatory sync mechanism is currently active, the countdown counter register \$ANS_DCOUNT will be updated in an automatic manner, such that no `sync acknowledgment` signal is left unaccounted for. The automatic non-participatory sync mechanism is implemented in dedicated hardware logic, preferably an instance of it in each tile 4, though in other embodiments it is not excluded that it could instead be implemented centrally for a group of tiles or all tiles.

With regard to the exit state behaviour, there are in fact a number of possibilities depending on implementation. In embodiments, to obtain the globally aggregated exit state, the sync logic 36, 76 only aggregates the local exit states from those tiles 4 in the specified sync group that executed a SYNC instruction, and not those that/those that executed a SANS instruction (the abstaining tile or tiles). Alternatively, the globally aggregated exit state is obtained by aggregating the local exit states from all the tiles 4 in the sync group that executed a SYNC and those that executed a SANS (both the participating and abstaining tiles 4). In the latter case, the local exit state output by the abstaining tile(s) 4 for global aggregation may be the actual locally aggregated exit state of that tile's workers at the time of executing the SANS, just as with the

SYNC instruction (see description of local consensus register \$LC 38). Alternatively the local “exit state” output by the abstaining tile 4 may be a default value, for instance the true value (e.g. logic 1) in embodiments where the exit state is binary. This prevents the abstaining tile 4 from interfering with the global exit state in embodiments where any false local exit state causes the global exit state to be false.

Regarding the return of the global exit state, there are two possibilities for this, regardless of whether or not the abstaining tile submits a local exit state for producing the global aggregate, and regardless of whether that value was an actual value or a default value. That is, in one implementation, the global aggregate exit state produced by the sync logic 36, 76 in the interconnect 34, 72 is stored only in the global consensus registers \$GC 42 of the participating tiles 4, which executed a SYNC instruction, and not the abstaining tiles 4 which instead executed a SANS instruction. In embodiments, instead a default value is stored in the global consensus register \$GX 42 of the tile(s) 4 that executed a SANS (the abstaining tiles). For instance this default value may be true, e.g. logic 1, in the case of a binary global exit state. In an alternative implementation however, the actual global aggregate produced by the sync logic 36, 76 is stored in the global consensus registers \$GC 42 of both the participating tiles 4 which executed SYNC instructions and the abstaining tiles 4 which instead executed a SANS instruction. Thus all tiles in the group may still have access to the globally aggregated exit state.

Figure 15 illustrates an example application of the processor architecture disclosed herein, namely an application to machine intelligence.

As will be familiar to a person skilled in the art of machine intelligence, machine intelligence begins with a learning stage where the machine intelligence algorithm learns a knowledge model. The model comprises a graph of interconnected nodes (i.e. vertices) 102 and edges (i.e. links) 104. Each node 102 in the graph has one or more input edges and one or more output edges. Some of the input edges of some of the nodes 102 are the output edges of some others of the nodes, thereby connecting together the nodes to form the graph. Further, one or more

of the input edges of one or more of the nodes 102 form the inputs to the graph as a whole, and one or more of the output edges of one or more of the nodes 102 form the outputs of the graph as a whole. Sometimes a given node may even have all of these: inputs to the graph, outputs from the graph and connections to other nodes. Each edge 104 communicates a value or more often a tensor (n-dimensional matrix), these forming the inputs and outputs provided to and from the nodes 102 on their input and output edges respectively.

Each node 102 represents a function of its one or more inputs as received on its input edge or edges, with the result of this function being the output(s) provided on the output edge or edges. Each function is parameterized by one or more respective parameters (sometimes referred to as weights, though they need not necessarily be multiplicative weights). In general the functions represented by the different nodes 102 may be different forms of function and/or may be parameterized by different parameters.

Further, each of the one or more parameters of each node's function is characterized by a respective error value. Moreover, a respective condition may be associated with the error(s) in the parameter(s) of each node 102. For a node 102 representing a function parameterized by a single parameter, the condition may be a simple threshold, i.e. the condition is satisfied if the error is within the specified threshold but not satisfied if the error is beyond the threshold. For a node 102 parameterized by more than one respective parameter, the condition for that node 102 having reached an acceptable level of error may be more complex. For example, the condition may be satisfied only if each of the parameters of that node 102 falls within respective threshold. As another example, a combined metric may be defined combining the errors in the different parameters for the same node 102, and the condition may be satisfied on condition that the value of the combined metric falls within a specified threshold, but otherwise the condition is not satisfied if the value of the combined metric is beyond the threshold (or vice versa depending on the definition of the metric). Whatever the condition, this gives a measure of whether the error in the parameter(s) of the node falls below a certain

level or degree of acceptability. In general any suitable metric may be used. The condition or metric may be the same for all nodes, or different for different respective ones of the nodes.

In the learning stage the algorithm receives experience data, i.e. multiple data points representing different possible combinations of inputs to the graph. As more and more experience data is received, the algorithm gradually tunes the parameters of the various nodes in the graph based on the experience data so as to try to minimize the errors in the parameters. The goal is to find values of the parameters such that the output of the graph is as close as possible to a desired output for a given input. As the graph as a whole tends toward such a state, the graph is said to converge. After a suitable degree of convergence the graph can then be used to perform predictions or inferences, i.e. to predict an outcome for some given input or infer a cause for some given output.

The learning stage can take a number of different possible forms. For instance, in a supervised approach, the input experience data takes the form of training data, i.e. inputs which correspond to known outputs. With each data point, the algorithm can tune the parameters such that the output more closely matches the known output for the given input. In the subsequent prediction stage, the graph can then be used to map an input query to an approximate predicted output (or vice versa if making an inference). Other approaches are also possible. For instance, in an unsupervised approach, there is no concept of a reference result per input datum, and instead the machine intelligence algorithm is left to identify its own structure in the output data. Or in a reinforcement approach, the algorithm tries out at least one possible output for each data point in the input experience data, and is told whether this output is positive or negative (and potentially a degree to which it is positive or negative), e.g. win or lose, or reward or punishment, or such like. Over many trials the algorithm can gradually tune the parameters of the graph to be able to predict inputs that will result in a positive outcome. The various approaches and algorithms for learning a graph will be known to a person skilled in the art of machine learning.

According to an exemplary application of the techniques disclosed herein, each worker thread is programmed to perform the computations associated with a respective individual one of the nodes 102 in a machine intelligence graph. In this case at least some of the edges 104 between nodes 102 correspond to the exchanges of data between threads, and some may involve exchanges between tiles. Furthermore, the individual exit states of the worker threads are used by the programmer to represent whether or not the respective node 102 has satisfied its respective condition for convergence of the parameter(s) of that node, i.e. has the error in the parameter or parameters fallen within the acceptable level or region in error space. For instance, this is one example use of the embodiments where each of the individual exit states is an individual bit and the aggregated exit state is an AND of the individual exit states (or equivalently an OR if 0 is taken to be positive); or where the aggregated exit state is a trinary value representing whether the individual exit states were all true, all false or mixed. Thus, by examining a single register value in the exit state register 38, the program can determine whether the graph as whole, or at least a sub-region of the graph, has converged to an acceptable degree.

As another variant of this, embodiments may be used where the aggregation takes the form of a statistical aggregation of individual confidence values. In this case each individual exit state represents a confidence (e.g. as a percentage) that the parameters of the node represented by the respective thread have reached an acceptable degree of error. The aggregated exit state can then be used to determine an overall degree of confidence as to whether the graph, or a subregion of the graph, has converged to an acceptable degree.

In the case of a multi-tile arrangement 6, each tile runs a subgraph of the graph. Each subgraph comprises a supervisor subprogram comprising one or more supervisor threads, and a set of worker threads in which some or all of the workers may take the form of codelets.

In such applications, or indeed any graph-based application where each worker thread is used to represent a respective node in a graph, the “codelet” comprised by each worker may be

defined as a software procedure operating on the persistent state and the inputs and/outputs of one vertex, wherein the codelet:

- is launched on one worker thread register context, to run in one barrel slot, by the supervisor thread executing a "run" instruction;
- runs to completion without communication with other codelets or the supervisor (except for the return to the supervisor when the codelet exits);
- has access to the persistent state of a vertex via a memory pointer provided by the "run" instruction, and to a non-persistent working area in memory which is private to that barrel slot; and
- executes "EXIT" as its last instruction, whereupon the barrel slot which it was using is returned to the supervisor, and the exit state specified by the exit instruction is aggregated with the local exit state of the tile which is visible to the supervisor.

To update a graph (or sub-graph) means to update each constituent vertex once, in any order consistent with the causality defined by the edges. To update a vertex means to run a codelet on the vertex state. A codelet is an update procedure for vertices – one codelet is usually associated with many vertices. The supervisor executes one RUN instruction per vertex, each such instruction specifying a vertex state address and a codelet address.

It will be appreciated that the above embodiments have been described by way of example only.

For instance, the applicability of the exit state aggregation mechanism is not limited to the above-described architecture in which a separate context is provided for the supervisor thread, or in which the supervisor thread runs in a slot and then relinquishes its slot to a worker. In another arrangement for example, the supervisor may run in its own dedicated slot.

Further, the terms “supervisor” and “worker” do not imply any specific responsibilities except where otherwise explicitly stated, and particularly do not in themselves necessarily limit to the

above-described scheme in which a supervisor thread relinquishes its time slot to a worker, and so forth. In general, worker thread may refer to any thread to which some computational task is allocated. The supervisor may represent any kind of overseeing or coordinating thread responsible for actions such as: assigning workers to barrel slots, and/or performing barrier synchronizations between multiple threads, and/or performing any control-flow operation (such as a branch) in dependence on the outcome of more than one thread.

Where reference is made to a sequence of interleaved time slots, or the like, this does not necessarily imply that the sequence referred to makes up all possible or available slots. For instance, the sequence in question could be all possible slots or only those currently active. It is not necessarily precluded that there may be other potential slots that are not currently included in the scheduled sequence.

The term tile as used herein does not necessarily limit to any particular topography or the like, and in general may refer to any modular unit of processing resource, comprising a processing unit 10 and corresponding memory 11, in an array of like modules, typically at least some of which are on the same chip (i.e. same die).

Further, the scope of the present disclosure is not limited to a time-deterministic internal interconnect or a non-time-deterministic external interconnect. The synchronization and aggregation mechanisms disclosed herein can also be used in a completely time-deterministic arrangement, or a completely non-time-deterministic arrangement.

Furthermore, where reference is made herein to performing a synchronization or an aggregation between a group of tiles, or a plurality of tiles or the like, this does not necessarily have to refer to all tile on the chip or all tiles in the system unless explicitly stated. E.g. the SYNC and EXIT instructions could be configured to perform the synchronization and aggregation only in relation to a certain subset of tiles 4 on a given chip and/or only a subset of chips 2 in a given system; whilst some other tiles 4 on a given chip, and/or some other chips in a given system,

may not be involved in a given BSP group, and could even be being used for some completely separate set of tasks unrelated to the computation being performed by the group at hand.

Also, while certain modes of SYNC instruction have been described above, the scope of the present disclosure more generally is not limited to such modes. For instance, the list of modes given above is not necessarily exhaustive. Or in other embodiments, the SYNC instruction may have fewer modes, e.g. the SYNC need not support different hierarchical levels of external sync, or need not distinguish between on-chip and inter-chip syncs (i.e. in an inter-tile mode, always acts in relation to all tiles regardless of whether on chip or off chip). In yet further alternative embodiments, the SYNC instruction need not take a mode as an operand at all. E.g. in embodiments separate versions of the SYNC instruction (different opcodes) may be provided for the different levels of synchronization and exit state aggregation (such as different SYNC instructions for on-tile synchronization and inter-tile, on-chip synchronization). Or in other embodiments, a dedicated SYNC instruction may only be provided for inter-tile synchronizations (leaving on-tile synchronization between threads, if required, to be performed in general purpose software).

Furthermore, the sync zones are not limited to being hierarchical (i.e. one nested in another), and in other embodiments the selectable sync zones may consist of or include one or more non-hierarchical groups (all tiles of that group not nested within a single other selectable group).

Further, the above-described synchronization schemes do not exclude the involvement, in embodiments, of external resources other than multi-tile processors, e.g. a CPU processor such as the host processor, or even one or more components that are not processors such as one or more network cards, storage devices and/or FPGAs. For instance, some tiles may elect to engage in data transfers with an external system where these transfers form the computational burden of that tile. In this case the transfers should be completed before the next barrier. In some cases the exit state of the tile may depend on a result of the communication with the

external resource, and this the resource may vicariously influence the exit state. Alternatively or additionally, resources other than multi-tile processors, e.g. the host or one or more FPGAs, could be incorporated into the synchronization network itself. That is to say, a sync signal such as a Sync_req is required from this/these additional resources in order for the barrier synchronization to be satisfied and the tiles to proceed to the next exchange phase. Further, in embodiments the aggregated global exit state may include in the aggregation an exit state of the external resource, e.g. from an FPGA.

Other applications and variants of the disclosed techniques may become apparent to a person skilled in the art once given the disclosure herein. The scope of the present disclosure is not limited by the described embodiments but only by the accompanying claims.

Claims

1. A processing system comprising an arrangement of tiles and an interconnect for communicating between the tiles, wherein:

each tile comprises an execution unit for executing machine code instructions, each being an instance of a predefined set of instruction types in an instruction set of the processing system, each instruction type in the instruction set being defined by a corresponding opcode and zero or more operand fields for taking zero or more operands;

the interconnect is operable to conduct communications between a group of some or all of the tiles according to a bulk synchronous parallel scheme, whereby each of the tiles in said group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all the tiles in said group have completed the compute phase, wherein each tile in the group has a local exit state upon completion of the compute phase;

the instruction set comprises a synchronization instruction for execution by each tile in the group upon completion of its compute phase, wherein execution of the synchronization instruction causes the execution unit to send a synchronization request to hardware logic in the interconnect; and

the logic in the interconnect is configured to aggregate the local exit states into a global exit state, and, in response to completion of the compute phase by all the tiles in the group as indicated by receiving the synchronization request from all the tiles in the group, to store the global exit state in a global exit state register on each of the tiles in the group, thereby making the global exit state accessible by a portion of code running on each of the tiles in the group.

2. The processing system of claim 1, wherein the execution unit on each tile is configured to pause instruction issue in response to executing the synchronization instruction; and wherein the logic in the interconnect is configured so as, in response to receiving the synchronization request from all the tiles in the group, to send back a synchronization acknowledgment signal to each of the tiles in the group in order to resume the instruction issue.

3. The processing system of claim 1 or claim 2, wherein each of the local exit states and the global exit states is a single bit.
4. The processing system of claim 3, wherein the aggregation consists of a Boolean AND of the local exit states or a Boolean OR of the local exit states.
5. The processing system of claim 1 or 2, wherein the aggregated exit state comprises at least two bits representing a trinary value, indicating whether the local exit states are all true, all false or mixed.
6. The processing system of any one of claims 1 to 5, wherein each of said group of tiles comprises a local exit state register arranged to represent the local exit state of the tile.
7. The processing system of any one of claims 1 to 6, wherein each tile in the group comprises:
 - multiple context register sets, each context register set each arranged to store a program state of a respective one of multiple threads; and
 - a scheduler arranged to schedule the execution of a respective one of a plurality of worker threads in each of a plurality of time slots in a repeating sequence of interleaved time slots, with the program state of each of the worker threads being stored in a respective one of said context register sets;
 - wherein according to said bulk synchronous parallel scheme, the exchange phase is held back until all the worker threads on all the tiles in the group have completed the compute phase;
 - wherein the local exit state on each tile is an aggregate of an individual exit state output by each of the worker threads on the tile; and
 - wherein said portion of code comprises at least one of the multiple threads on the tile.
8. The processing system of claim 7, wherein each tile in the group comprises hardware logic configured to perform said aggregation of the individual exit states into the local exit state.

9. The processing system of claim 8, wherein the instruction set comprises an exit instruction for inclusion in each of the worker threads, the execution unit being configured to output the individual exit state of the respective worker thread and terminate the respective worker thread in response to the opcode of the exit instruction.

10. The processing system of claim 7, 8 or 9, wherein each of the individual exit states and the local exit states is a single bit, and the aggregation of the individual exit states consists of a Boolean AND of the individual exit states or a Boolean OR of the individual exit states.

11. The processing system of claim 7, 8 or 9, wherein the local exit state comprises at least two bits representing a trinary value, indicating whether the individual exit states are all true, all false, or mixed.

12. The processing system of any of claims 7 to 11, wherein the exchange phase is arranged to be performed by a supervisor thread separate to the worker threads, and said at least one thread comprises the supervisor thread.

13. The processing system of claim 12 when dependent on claim 2, wherein the pausing of the instruction issue comprises at least pausing issue of instructions from the supervisor thread pending the synchronization acknowledgment.

14. The processing system of claim 12 or 13, wherein the context register sets on each tile comprise: multiple worker context register sets arranged to represent the program state of respective ones of said plurality of worker threads, and an additional supervisor context register set comprising an additional set of registers arranged to represent a program state of the supervisor thread.

15. The processing system of claim 14, wherein:
the supervisor thread is arranged to begin by running in each of said time slots;
the instruction set further comprises a relinquish instruction and the execution unit is configured so as, in response to the opcode of the relinquish instruction, to relinquish the

time slot in which the relinquish instruction is executed to the respective worker thread;
and

the exit instruction causes the respective time slot in which the exit instruction is executed to be handed back to the supervisor thread, such that the supervisor thread resumes running in the respective slot.

16. The processing system of any one of claims 1 to 15, programmed with said code; wherein said portion of code is configured to use the global exit state, once valid, to perform a branch decision which depends on the global exit state.

17. The processing system of any one of claims 1 to 16, programmed to perform a machine intelligence algorithm in which each node in a graph has one or more respective input edges and one or more respective output edges with the input edges of at least some of the nodes being the output edges of at least some others of the nodes, each node comprising a respective function relating its output edges to its input edges, with each respective function being parameterized by one or more respective parameters, and each of the respective parameters having an associated error, such that the graph converges toward a solution as the errors in some or all of the parameters reduce;

wherein each of the tiles models a respective subgraph comprising a subset of the nodes in the graph, and each of the local exit states is used to indicate whether the errors in the one or more parameters of the nodes in the respective subgraph have satisfied a predetermined condition.

18. The processing system of any one of claims 1 to 17, wherein said group is selected at least in part by an operand of the synchronization instruction.

19. The processing system of claim 18, wherein the operand of the synchronization instructions selects whether to include only tiles on the same chip or tiles on different chips in said group.

20. The processing system of claim 18 or 19, wherein the operand of the synchronization instructions selects said group from amongst different hierarchical levels of groupings.

21. The processing system of any one of claims 1 to 20, wherein the instruction set further includes an abstain instruction, which causes the tile on which the abstain instruction is executed to opt-out of said group.

22. A method of operating a processing system comprising an arrangement of tiles and an interconnect for communicating between the tiles, wherein each tile comprises an execution unit for executing machine code instructions, each being an instance of a predefined set of instruction types in an instruction set of the processor, each instruction type in the instruction set being defined by a corresponding opcode and zero or more operand fields for taking zero or more operands; the method comprising:

conducting communications between a group of some or all of the tiles, via the interconnect, according to a bulk synchronous parallel scheme, whereby each of the tiles in said group performs an on-tile compute phase followed by an inter-tile exchange phase with the exchange phase being held back until all the tiles in said group have completed the compute phase, wherein each tile in the group has a local exit state upon completion of the compute phase;

wherein the instruction set comprises a synchronization instruction for execution by each tile in the group upon completion of its compute phase, wherein execution of the synchronization instruction causes the execution unit to send a synchronization request to hardware logic in the interconnect; and

the method comprises, in response to completion of the compute phase by all the tiles in the group as indicated by receiving the synchronization request from all the tiles in the group, triggering the logic in the interconnect to aggregate the local exit states into a global exit state, and to store the global exit state in a global exit state register on each of the tiles in the group, thereby making the global exit state accessible by a portion of code running on each of the tiles in the group.

23. A computer program product comprising code embodied on computer readable storage and configured to execute on the processing system of any of claims 1 to 21, the code comprising a portion for execution on each tile in the group including an instance of the synchronization instruction in each portion.



Figure 1**Figure 2**



Figure 3Figure 4


Figure 5


Figure 6



Figure 7**Figure 8**


Figure 9

Figure 10





Figure 11


Figure 12


Figure 13

Figure 14





Figure 15

Figure 16

