



- (51) **International Patent Classification:**
G06F 9/50 (2006.01) *G06F 11/14* (2006.01)
- (21) **International Application Number:**
PCT/US2017/039760
- (22) **International Filing Date:**
28 June 2017 (28.06.2017)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
15/202,044 05 July 2016 (05.07.2016) US
- (71) **Applicants:** SAUDI ARABIAN OIL COMPANY [SA/SA]; 1 Eastern Avenue, Dhahran, 31311 (SA). KING FAHD UNIVERSITY OF PETROLEUM AND MINERALS [SA/SA]; P. O. Box 5041, Dhahran, 31261 (SA).
- (71) **Applicant (for AG only):** ARAMCO SERVICES COMPANY [US/US]; 9009 West Loop South, Houston, Texas 77096 (US).
- (72) **Inventors:** ALSHAIKH, Raed Abdullah; P.O. Box 5382, Dhahran, 31311 (SA). SAIT, Sadiq M.; P. O. Box 673, Dhahran, 31261 (SA).
- (74) **Agent:** RHEBERGEN, Constance G. et al.; Bracewell LLP, P.O. Box 61389, Houston, Texas 77208-1389 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) **Title:** HIGH PERFORMANCE AND GRID COMPUTING WITH LIVELINESS AND DEADLINES FAULT TOLERANT DATA DISTRIBUTOR QUALITY OF SERVICE





FIG. 2

(57) **Abstract:** High performance computing (HPC) and grid computing processing for seismic and reservoir simulation are performed without impacting or losing processing time in case of failures. A Data Distribution Service (DDS) standard is implemented in High Performance Computing (HPC) and grid computing platforms, to avoid the shortcomings of current Message Passing Interface (MPI) communication between computing modules, and provide quality of service (QoS) for such applications. QoS properties of the processing can be controlled. Multiple data publishers or master nodes of a cluster have access to the same data source. Each of these publishers has an ownership strength quality of service, and the publisher with the highest ownership strength number is the designated publisher of the data to subscriber processor nodes of the cluster. If the designated data publisher prematurely terminates or crashes for some reason, then the publisher node with the next highest ownership strength measure is designated as data publisher and continues publishing data to subscribers. The QoS properties include ability to switch to a different designated publisher when the input data is being processed in real time, and a specified deadline time has passed during which data has not been published.



SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

HIGH PERFORMANCE AND GRID COMPUTING WITH LIVELINESS AND DEADLINES FAULT TOLERANT DATA DISTRIBUTOR QUALITY OF SERVICE

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The present invention relates to high performance and grid computing of data for exploration and production of hydrocarbons, such as computerized simulation of hydrocarbon reservoirs in the earth, geological modeling, and processing of seismic survey data, and in particular to quality of service (QoS) control of such computing.

2. Description of the Related Art

[0002] In the oil and gas industries, massive amounts of data are required to be processed for computerized simulation, modeling and analysis for exploration and production purposes. For example, the development of underground hydrocarbon reservoirs typically includes development and analysis of computer simulation models of the reservoir. These underground hydrocarbon reservoirs are typically complex rock formations which contain both a petroleum fluid mixture and water. The reservoir fluid content usually exists in two or more fluid phases. The petroleum mixture in reservoir fluids is produced by wells drilled into and completed in these rock formations.

[0003] A geologically realistic model of the reservoir, and the presence of its fluids, also helps in forecasting the optimal future oil and gas recovery from hydrocarbon reservoirs. Oil and gas companies have come to depend on geological models as an important tool to enhance the ability to exploit a petroleum reserve. Geological models of reservoirs and oil/gas fields have become increasingly large and complex.

[0004] In simulation and geological models, the reservoir is organized into a number of individual cells. Seismic data with increasing accuracy has permitted the cells to be on the order of 25 meters areal (x and y axis) intervals. For what are known as giant reservoirs, the number of cells is the least hundreds of millions, and reservoirs of what is known as giga-cell size (a billion cells or more) are encountered.

[0005] Similar considerations of data volume are also presented in seismic data processing. Seismic data obtained from surveys over large areas of the earth's surface such as above giant reservoirs, has been acquired and made available in increased volumes. In processing vast amounts of data of all three of the types described above, processing time was an important consideration.

[0006] Three types of computer systems have been available for processing the vast amounts of data of the types encountered in petroleum exploration and production. These are supercomputers, high performance computing (HPC) and grid computing. Typically, supercomputers are specially designed for particular calculation intensive tasks. An HPC system takes the form of a group of powerful workstations or servers, joined together as a network to function as one supercomputer. Grid computing involves a more loosely coupled, heterogeneous and often dispersed network of workstations or servers than HPC.

[0007] So far as is known, existing distributed memory HPC and grid computing systems did not provide proper quality of service (QoS) based communication because of two limitations. First, standard communication libraries such as Message Passing Interface (MPI) and Parallel Virtual Machine (PVM) did not provide a capability for applications to specify service quality for computation and communication. Second, modern high-speed interconnects such as InfiniBand, Myrinet, Quadrics and Gigabit Ethernet were optimized for performance rather than for predictability of communication latency and bandwidth.

[0008] There has been, so far as is known, little attention given to QoS control in high performance and grid computing. HPC users have witnessed a dramatic increase in performance over the last ten years with regard to the HPC systems. What used to take one month of HPC computation time in the ten years ago, is now taking only a few hours to run in current systems.

[0009] In view of this, the simplest remedy to users for a data accuracy failure rate during a routine computation run has been resubmitting the processing data run after disregarding or offlining (or what is known as fencing) the problematic node/core. However, the hours of the crashed job were thus discarded and wasted. Moreover, in the case of a more extensive data processing run which would require several days or even weeks to perform, resubmitting the entire data set for processing was required. This was duplicative and time consuming.

[0010] U. S. Patent No. 7,526,418, which is owned by the assignee of the present application, relates to a simulator for giant hydrocarbon reservoirs composed of a massive number of cells. The simulator mainly used high performance computers (HPC). Communication between the cluster computers was performed according to conventional, standard methods, such as MPI mentioned above and Open MP.

[0011] U. S. Published Patent Application No. 2011/0138396 related to a data distribution mechanism in HPC clusters. The focus of the system described was methodology to enable data to be distributed rapidly to various computation nodes in an HPC cluster. Thus, the focus and teachings of this system were improving processing speed by more rapidly distributing data to the cluster nodes.

SUMMARY OF THE INVENTION

[0012] Briefly, the present invention provides a new and improved computer implemented method of computerized processing in a data processing system of data for exploration and production of hydrocarbons. The data processing system includes a plurality of master nodes, each with an established quality of service standard profile. The quality of service profile includes assigned ownership strength, with the master node of the plurality of master nodes having a highest assigned ownership strength being designated master publisher node for the exploration and production data being processed. The quality of service profile further includes a specified deadline time interval for periodic production of data by the designated master publisher node when the data being processed is being produced in real time. The data processing system also includes a plurality of processor nodes established as subscribers to receive exploration and production data from the designated master publisher, and a data memory. According to the computer implemented method, the established quality of service standard profile is transmitted from the designated master publisher to the subscriber processor nodes. The designated master publisher establishes a domain for exploration and production processing by the designated master publisher and designated ones of the pluralities of processor nodes as subscriber processor nodes. The exploration and production data is produced from the designated master publisher to the designated subscriber processor nodes of the domain, and a determination made if the data being processed is indicated to be real time data. The transmitted exploration and production data are processed in the designated subscriber processor nodes of the domain, and the processed exploration and production data of the designated subscriber processor nodes is monitored at the designated master publisher of the domain. The processed exploration and production data received at the designated master publisher is assembled in the data memory of the data processing system. Performance of the designated master publisher node is monitored during

the processing of the exploration and production data to determine if the specified time interval has passed without periodic production of data from the designated master publisher node. If the specified time interval has not passed, monitoring of performance of the designated master publisher node is continued to determine that the designated master publisher node is operating and processing of the exploration and production data is continued. If the designated master publisher node is not operating, the ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service profile is determined. A master node having the highest ownership strength quality of service profile is then established as the designated master publisher node the processing of the exploration and production data is resumed.

[0013] The present invention also provides a new and improved data processing system for computerized processing of data for exploration and production of hydrocarbons. The data processing system includes a plurality of master nodes, each with an established quality of service standard profile including assigned ownership strength. The master node of the plurality of master nodes having highest assigned ownership strength is designated master publisher node for the exploration and production data being processed. The quality of service profile also includes a specified deadline time interval for periodic production of data by the designated master publisher node when the data being processed is being produced in real time. The data processing system also includes a plurality of processor nodes and a data memory. In the data processing system, the designated master publisher node transmits the established quality of service standard profile from the designated master publisher to the subscriber processor nodes. The designated master publisher node also establishes a domain for exploration and production processing by the designated master publisher and designated ones of the plurality of processor nodes as subscriber processor nodes, and produces the exploration and production data to the designated subscriber processor nodes of the domain.

The master node also determines if the data being processed is indicated to be real time data. The processor nodes established as subscribers receive exploration and production data from the designated master publisher, and receive in the designated subscriber processor nodes of the domain the exploration and production data and the established quality of service standard profile from the designated master publisher node. The designated subscriber processor nodes of the domain process the transmitted exploration and production data. The master node further monitors the processed exploration and production data of the designated subscriber processor nodes of the domain, and assembles in the data memory the processed exploration and production data from the designated subscriber processor nodes which comply with the transmitted established quality of service standard profile. The master node also monitors performance of the designated master publisher node during the processing of the exploration and production data to determine if the specified time interval has passed without periodic production of data from the designated master publisher node. If the specified time interval has not passed, monitoring performance of the designated master publisher node to determine that the designated master publisher node is operating, processing of the exploration and production data are continued. If the designated master publisher node is not operating, ownership strength of the others of the plurality of master nodes is determined according to their respective ownership strength quality of service profile. The master node having the highest ownership strength quality of service profile is established as designated master publisher node and the processing of the exploration and production data is resumed.

[0014] The present invention further provides a new and improved data storage device having stored in a computer readable medium computer operable instructions for causing a data processing system to perform computerized processing of data for exploration and production of hydrocarbons. The data processing system includes a plurality of master

nodes, each with an established quality of service standard profile including an assigned ownership strength, and with the master node of the plurality of master nodes having a highest assigned ownership strength being designated master publisher node for the exploration and production data being processed. The quality of service profile further includes a specified deadline time interval for periodic production of data by the designated master publisher node when the data being processed is being produced in real time. The data processing system includes a plurality of processor nodes established as subscribers to receive the exploration and production data from the designated master publisher node, and a data memory. The instructions stored in the data storage device causing the data processing system to transmit the established quality of service standard profile from the designated master publisher to the subscriber processor nodes, and to establishing with the designated master publisher a domain for exploration and production processing by the designated master publisher and designated ones of the plurality of processor nodes as subscriber processor nodes. The instructions also cause production of the exploration and production data from the designated master publisher to the designated subscriber processor nodes of the domain. The instructions also cause a determination to be made if the data being processed is indicated to be real time data. The instructions cause processing of the transmitted exploration and production data in the designated subscriber processor nodes of the domain, and monitoring at the designated master publisher the processed exploration and production data of the designated subscriber processor nodes of the domain. The instructions also cause assembling in the data memory of the data processing system the processed exploration and production data received at the designated master publisher. The instructions further cause monitoring of performance of the designated master publisher node during the processing of the exploration and production data to determine if the specified time interval has passed without periodic production of data from the designated master

publisher node. If the specified time interval has not passed, the instructions cause continuing the step of monitoring performance of the designated master publisher node to determine that the designated master publisher node is operating, and continuing the processing of the exploration and production data. If the designated master publisher node is not operating, the instructions cause ownership strength of the others of the plurality of master nodes to be determined according to their respective ownership strength quality of service profile, and establishing the master node having the highest ownership strength quality of service profile as designated master publisher node, and resuming the processing of the exploration and production data.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] Figure 1 is a schematic block diagram of a prior art data processing system for high performance computing.

[0016] Figure 2 is a schematic block diagram of a data processing system for high performance and grid computing with quality of service control according to the present invention.

[0017] Figure 3 is a functional block diagram of the data processing system of Figure 2 configured for high performance processing with quality of service control according to the present invention.

[0018] Figure 4 is a functional block diagram of a set of data processing steps performed in the data processing system of Figures 2 and 3 for high performance processing with quality of service control according to the present invention.

[0019] Figure 5 is a functional block diagram of a set of data processing steps performed in the data processing system of Figures 2 and 3 for high performance processing with quality of service control according to the present invention.

[0020] Figure 6 is a functional block diagram indicating the functional interrelationship between the interactive operations for quality of service control according to the present invention.

[0021] Figures 6A, 6B, 6C and 6D are functional block diagrams indicating in more detail the interactive operation of processors of the data processing system of Figures 2 and 3 during performance of the data processing steps of Figures 4 and 5.

[0022] Figure 7 is a functional block diagram indicating the operation of processors of the data processing system of Figures 2 and 3 according to the ownership, deadlines and liveness strength quality of service controls according to the present invention.

[0023] Figure 8 is a plot of runtime for the data processing system shown in Figures 2 and 3 for high performance processing with quality of service control with different sizes of input data in comparison with prior art MPI communication protocols.

[0024] Figure 9 is a plot of network time delay in engaging a new processor node for different file sizes in the data processing system of Figures 2 and 3.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0025] The present invention relates to high performance and grid computing of data for exploration and production of hydrocarbons, such as computerized simulation of hydrocarbon reservoirs in the earth, geological modeling, processing of seismic survey data, and other types of data gathered and processed to aid in the exploration and production of

hydrocarbons. For the purposes of the present invention data of the foregoing types are referred to herein as exploration and production data. The present invention is particularly adapted for processing exploration and production data where vast amounts of such data are present, such as in or around what are known as giant reservoirs.

[0026] In the drawings, Figure 1 represents an example prior art high performance computing network P. The high performance computing network P is configured for parallel computing using the message passing interface (MPI) with a master node 10 transferring data through what are known as serial heartbeat connections over data links 12 of a management network 14 to a number of processor nodes 16. The processor nodes 16 are configured to communicate with each other as indicated at 18 according to the message passing interface (MPI) standard communication library during parallel computing and processing of data. As has been set forth, so far as is known, standard communication libraries such as Message Passing Interface (MPI) and Parallel Virtual Machine (PVM) did not provide a capability for applications to specify service quality for computation and communication.

[0027] With the present invention, as is shown schematically in Figure 2 in a data processing system D a plurality of master nodes 20 of a CPU 22 and a group of processor or worker nodes 24 operating as a network arranged for high performance or grid computing, depending on the configuration of the network, of exploration and production data. As will be set forth, the data processing system D processes exploration and production data with a controllable specified quality of service (QoS) for the processing applications. Data processing system D operates according to the processing techniques which are shown schematically in Figures 4, 5, 6A through 6D, and 7. Thus, high performance computing (HPC) and grid computing processing of exploration and production data are performed without impacting or losing processing time in case of failures. A data distribution service

(DDS) standard is implemented in the high performance computing (HPC) and grid computing platforms of the data processing system D, to avoid shortcomings of message passing interface (MPI) communication between computing modules, and provide quality of service (QoS) for such applications.

[0028] Considering now the data processing system according to the present invention, as illustrated in Figure 2, the data processing system D is provided as a processing platform for high performance computing (HPC) and grid computing of exploration and processing data. The data processing system D includes one or more central processing units or CPU's 22. The CPU or CPU's 22 have associated therewith a reservoir memory or database 26 for general input parameters, of a type and nature according to the exploration and production data being processed, whether reservoir simulation, geological modeling, seismic data or the like.

[0029] A user interface 28 operably connected with the CPU 22 includes a graphical display 30 for displaying graphical images, a printer or other suitable image forming mechanism and a user input device 32 to provide a user access to manipulate, access and provide output forms of processing results, database records and other information.

[0030] The reservoir memory or database 26 is typically in a memory 34 of an external data storage server or computer 38. The reservoir database 26 contains data including the structure, location and organization of the cells in the reservoir model, data general input parameters, as well as the exploration and production data to be processed, as will be described below.

[0031] The CPU or computer 22 of data processing system D includes the master nodes 20 and an internal memory 40 coupled to the master nodes 20 to store operating instructions, control information and to serve as storage or transfer buffers as required. The data

processing system D includes program code 42 stored in memory 40. The program code 42, according to the present invention, is in the form of computer operable instructions causing the master nodes 20 and processor nodes 24 to transfer the exploration and production data and control instructions back and forth according to data distribution service (DDS) intercommunication techniques, as will be set forth.

[0032] It should be noted that program code 42 may be in the form of microcode, programs, routines, or symbolic computer operable languages that provide a specific set of ordered operations that control the functioning of the data processing system D and direct its operation. The instructions of program code 42 may be stored in memory 40 or on computer diskette, magnetic tape, conventional hard disk drive, electronic read-only memory, optical storage device, or other appropriate data storage device having a computer usable medium stored thereon. Program code 42 may also be contained on a data storage device as a computer readable medium.

[0033] The processor nodes 24 are general purpose, programmable data processing units programmed to perform the processing of exploration and production data according to the present invention. The processor nodes 24 operate under control of the master node 20 and the processing results obtained are then assembled in memory 34 where the data are provided for formation with user interface 28 of output displays to form data records for analysis and interpretation.

[0034] Although the present invention is independent of the specific computer hardware used, an example embodiment of the present invention is preferably based on master nodes 20 and processor nodes 24 of an HP Linux cluster computer. It should be understood, however, that other computer hardware may also be used.

[0035] According to the present invention, the data processing system D whose components are shown in Figure 2 is configured as shown in Figure 3 according to the data distribution service (DDS) techniques. As indicated in Figure 3, a designated one of the master nodes 20 also shown in Figure 2 is established as a designated publisher master node as indicated at 50 in that such master node 20 is the node responsible for dissemination and distribution of the exploration and production data to be processed by the processor nodes 24. The designated publisher master node 20 includes a persistence service capability as shown at 51 to preserve data samples so that they can be furnished to processor nodes which replace failed processor nodes, as will be described. The designated publisher master node 20 is also as indicated at 52 in Figure 3 established as a data writer for the exploration and production data distributed for processing, so that values of the data which is distributed can be established.

[0036] The processor nodes 24 in operating according to data distribution service (DDS) are each individually established as subscribers, as indicated at 54 in Figure 3. Thus the processor nodes as subscribers are configured to operate as the processors responsible assigned to receive the exploration and production data received as a result of the subscriber relationship to the publisher 50 of the master node 20. The processor nodes 24 are also configured as data readers as indicated at 56 in Figure 3. As data readers 56, the processor nodes 24 receive an allocated portion of the exploration and production data from the data writer 52 of the master node 20.

[0037] The DDS techniques of the present invention are further explained in the data distribution service (DDS) standard. The DDS standard provides a scalable, platform-independent, and location-independent middleware infrastructure to connect information producers to consumers (i.e. nodes to nodes). DDS also supports many quality-of-service

(QoS) policies, such as asynchronous, loosely-coupled, time-sensitive and reliable data distribution at multiple layers (e.g., middleware, operating system, and network).

[0038] The DDS methodology implements a publish/subscribe (PS) model for sending and receiving data, events, and commands among the participant or designated master node 20 and processor nodes 24 in the processing of exploration and production data according to the present invention. As will be set forth, the master node serves as a primary publisher and the processor nodes 24 function as the primary subscribers. The processor nodes 24 are also configured to transfer the processed exploration and production data results to the designated master node 20, and the designated master node 20 is configured for this purpose as a subscriber function. Thus each node of the data processing system D can serve as a publisher, subscriber, or both simultaneously.

[0039] Nodes of the data processing system D that are producing information (publishers) create "topics" which are parameters of interest for the data being processed, depending on the type of exploration and production data being processed. The nodes operating in the DDS mode take care of delivering the data samples to those subscribers that indicate an interest in that topic. In a preferred computer network according to the present invention, example implementations of DDS provide low latency messaging (as fast as 65 microseconds between nodes) and high throughput (up to 950Mbps).

[0040] As has been set forth, the present invention processes exploration and production data for giant reservoirs where vast amounts of data need to be processed. An example type of processing performed is two-way data streaming between master and processor nodes. Examples of data streaming usage in exploration and production data for giant reservoirs are United States Patent Nos. 7,596,480; 7,620,534; 7,660,711; 7,526,418; and 7,809,537. It should be understood that the present invention may also be used in connection with

communication between nodes for HPC or grid computing of exploration and production data for other types of processing in addition to data streaming. The data communication according to the present invention between nodes for HPC and grid computing may also be used for other types of data, such as bioinformatics processing and computational fluid dynamics. The present invention is adapted for use in processing of large amounts of data which consume considerable time and thus has an increased likelihood of system failure during the course of processing.

[0041] Figure 4 is a functional block diagram of a set 60 of data processing steps performed by the master node 20 in the data processing system D according to the present invention. As shown at step 62 of Figure 4, the master node 20 is designated as the master publisher. Master node 20 then spawns two threads using OpenMP to parallelize two functionalities. In the first thread, the master node 20 initializes during step 64 to be a publisher (P0) with selected QoS profile, which is predefined in an XML file. For example, in the data streaming embodiment described below, three main QoS policies are adopted. These are: durability, reliability, and history.

[0042] The “durability” of the QoS profile saves the published data topics so that the data topics can be delivered to subscribing nodes that join the system at a later time, even if the publishing node has already terminated. The persistence service can use a file system or a relational data base to save the status of the system.

[0043] The second QoS profile or policy, which is reliability, indicates the level of reliability requested by a DataReader or offered by a DataWriter. Publisher nodes may offer levels of reliability, parameterized by the number of past issues they can store for the purpose of retrying transmissions. Subscriber nodes may then request different levels of reliable delivery, ranging from fast-but-unreliable “best effort” to highly reliable in-order delivery.

Thus providing per data stream reliability control. In case the reliability type is set to “RELIABLE”, the write operation on the DataWriter may be blocked if the modification would cause data to be lost or else cause one of the resource limited to be exceeded.

[0044] The third policy, history, controls the behavior of the communication when the value of a topic changes before it is finally communicated to some of its existing DataReader entities. If the type is set to “KEEP_LAST”, then the service will only attempt to keep the latest values of the topic and discard the older ones. In this case, a specified value of depth of data retention regulates the maximum number of values the service will maintain and deliver. The default (and most common setting) for depth is one, indicating that only the most recent value should be delivered.

[0045] If the history type is set to “KEEP_ALL”, then the service will attempt to maintain and deliver all the values of the sent data to existing subscribers. The resources that the service can use to keep this history are limited by the settings of the RESOURCE_LIMITS QoS. If the limit is reached, then the behavior of the service will depend on the RELIABILITY QoS. If the reliability kind is “BEST_EFFORTS”, then the old values will be discarded. If the reliability setting is “RELIABLE”, then the service will block the DataWriter until it can deliver the necessary old values to all subscribers.

[0046] In the data streaming embodiment described herein, “DURABILITY” was used and specified, because it was desirable for compute nodes to continue joining the system whenever there is a failure in another node, and thus avoid the consequences of a node failure during an extended processing run. In the third policy of history, “KEEP_LAST” was selected since the streamed seismic or simulation data are not expected to change. Specifically, the seismic shots and simulation cell values are fixed and not dynamic. In reliability QoS policy, RELIABLE, was specified as all the values need to reach the

subscribers and have complete answers of the data streaming. The requirement of data accuracy prohibited missing or losing elements while transmitting.

[0047] In high performance and grid computing, multiple data publishers (master nodes) can exist in a cluster in which all of the master nodes function as publishers and have access to the same data source. With the present invention, a methodology is provided to enable fault tolerant continuity of the high performance and grid computing. According to the present invention, the quality of service (QoS) profile includes an OWNERSHIP_STRENGTH quality of service. Each of the data publisher master nodes has an assigned or delegated “OWNERSHIP_STRENGTH” quality of service, where the publisher with the highest ownership strength number is the designated publisher of the data. By default, OWNERSHIP_STRENGTH is assigned based on the time the master node is joining the domain (that is, the cluster). In other words, the first node joining the domain would be assigned the highest OWNERSHIP_STRENGTH integer. This default behavior can be manually changed by a user who may alter the OWNERSHIP_STRENGTH for a specific node (set it higher or lower) if it is preferred to re-arrange the priorities of the master nodes.

[0048] There are types of exploration and production data which require that the high performance cluster of data processing system D perform applications with real-time processing. In these situations, a periodic data stream is a main component as an input to the cluster. Examples of such data and applications include GeoSteering applications where the stream of data is processed real time in the cluster to provide processed output to steer drilling bits in real time drilling. Another example is real time reservoir simulation cluster where the processing of data happens in real time and simulation data input observed at periodic time intervals.

[0049] In real time processing applications, the data subscribers (compute nodes) of the cluster require and are programmed to expect that the stream of data be received in specific “periodical” time from the master node (Data publishers). This quality of service according to the present invention referred to as a “Deadline” QoS. With the Deadline quality of service or QoS, the present invention provides a capacity for cases where a specified deadline time is passed and data is not received. In such an event, the currently designated Data publisher master node is considered offline, and another data writer (master node) is chosen based on its Ownership Strength quality of service ranking under the Ownership Strength QoS functionality.

[0050] However, with the present invention, if the application does not need to be real time processing, the Deadline quality of service is neglected. In such cases, if the data publisher master node crashes, the compute nodes are not required to wait for data to be processed. The present invention still provides fault tolerant QoS, wherein a LIVELINESS quality of service determines the next data publisher to take the function of distributor master node based on its OWNERSHIP_STRENGTH quality of service, as discussed above.

[0051] If a master node currently serving as designated publisher master node and functioning as data publisher prematurely terminates or crashes for any reason, then according to the present invention, a substitute or replacement publisher master node with the next highest ownership strength is designated as master publisher node. The replacement publisher master node is then in charge of processing to continue publishing data to subscribers. Thus, the present invention provides fault tolerance between multiple data publishers. Detection of the active master node being prematurely terminated is accomplished as shown in Figure 7, as will be described below.

[0052] The present invention also provides a “LIVELINESS” or operating status inquiry quality of service profile. According to established selected time intervals, an inquiry is made of the operating status of the master node 20 presently serving as designated publisher master node. The inquiry is made in order to confirm the availability or presently active operating status of the currently serving designated master publisher node. The inquiry happens in the form of “network heart-beats” from the Data Distribution Service (DDS) to data writer (publisher) that has the strongest (OWNERSHIP) integer. If this publisher does not acknowledge these network heart-beat packets, then the LIVELINESS QoS assigns the next available master node that has the highest OWNERSHIP integer to be the publisher.

[0053] As indicated at step 66, the thread also specifies the domain where all the publishers and subscribers would work on, which is domain-0 in the present embodiment. Next, as indicated at step 68, a topic with an identifying name is specified and a DataWriter (DW-0) is initialized during step 70 under P0 using that topic. After that, the master node 20 as publisher begins reading the data matrices from input to be processed, and initializes the data structure for the source sample (SS) by defining the matrices dimension and the number of processor nodes 24 which are designated as processor or worker nodes for the processing of the exploration and production data. The master node then during step 72 starts sending the Source sample SS-0 to the designated worker processor nodes 24 through the DataWriter DW-0.

[0054] The second thread of master node 20 reverses the function of thread 0 by creating an instance of a subscriber S0 as indicated at step 74 with the selected QoS profile in Domain-0, in preparation to receive the partial results from the worker nodes 24 (designated worker nodes 24 act as subscribers at the beginning and then as publishers at the end of their processing). The master node 20 then listens to the worker nodes 24 to receive processing

results as indicated at 76 through the receiving sample RS-0 and verifies compliance with the selected QoS profile as indicated at step 78. The master node 20 checks the QoS profile while it is receiving data from the worker nodes. It only receives data from those working nodes 24 which are matching in their QoS. It will not receive and will not negotiate with other worker nodes which have incompatible QoS settings. The master node 20 then during step 80 stores the verified processing results in data memory, and the stored results are available for output and display.

[0055] Figure 5 is a functional block diagram of a set 84 of data processing steps performed by the processor nodes 24 in the data processing system D according to the present invention. On the subscribers' side (i.e., the workers), each node 24 during step 86 initiates itself as a subscriber to the main publisher P0, assigns an ID to itself (Wi), and during step 88 starts receiving the data for computational processing. The distribution of which data goes to which node 24 is done dynamically in a way that is determined by first identifying the data range taken by each node according to the format of the data. As indicated at step 90, the QoS checking is done during the nodes' receiving process. The data is then processed during step 92 according to the processing required. Each worker node 24 during step 94 then sends its output with verified QoS through its DataWriter (DW-i) to the master node 20 for result collection. As indicated at step 96, QoS checking is done during the sending of data in step 94. Thus, when there is communication between the publisher and subscriber, QoS checking is done to establish this connection.

[0056] Figure 6 is a functional block diagram indicating the functional interrelationship between the interactive operations shown in more detail in Figures 6A through 6D, respectively, for quality of service control according to the present invention. As shown in Figure 6, Figure 6B is a functional block diagram 100 illustrating the interaction of the data

processing nodes according to the present invention during the processing of exploration and production data. Figure 6A is a functional block diagram 150 of an initial set-up methodology performed prior to the actions shown in diagram 100. Figures 6C and Figure 6D are functional block diagrams of alternative control of the interactions of the data processing depending on the Quality of Service profile which is indicated to be in effect during the initial set-up.

[0057] Thus, as shown in Figure 6A, prior to step 102, the initial set-up methodology 150 is performed. The data processing application is submitted to the HPC cluster of the data processing system D as shown at 150 in Figure 6A. As indicated at step 152, a determination is made whether the data processing application is identified as one requiring real time processing. If real time processing is indicated during step 152, as indicated at step 154, the above described Deadline Quality of Service is established between the master node(s) 20 as published and the processor nodes 24. For the Deadline Quality of Service, a specified time is identified within which the currently designated publisher master node 20 must publish data to the worker nodes 24.

[0058] If it is determined during step 152 that real time processing is not indicated, then the Liveliness Quality of Service is established during step 156 between the publisher master node(s) 20 and the processor nodes 24. As a result of either of steps 154 or 156, processing proceeds to the step 128 (Figures 6A and 7) for selection of one of the nodes 20 as designated master node and designated master publisher as a result of having the highest ranking Ownership Strength QoS profile.

[0059] Figure 6B is a diagram 100 illustrating the interaction of the master node 20 selected during step 12B (Fig. 6A and 7) and worker processor nodes 24 in an example implementation of a data streaming processing of exploration and production data. The

implementation starts at step 102 by designating the master node 20 of the cluster as the main publisher.

[0060] After being so designated, the master node 20, in turn, spawns two threads using OpenMP to parallelize its two main functions: initializing the node to be a publisher (P0) with the selected QoS profile; and specifying during step 104 the domain which the publishers and subscribers are to work on, which is domain-0 in the example implementation. Specifying the domain is necessary in order to allow multiple groups of publishers and subscribers to work independently, segmenting the cluster into several smaller sub-clusters, if needed. Different algorithms may require different topics (i.e. datasets) to be sent independently by the same publisher, and each of these topics may have several DataWriters for redundancy.

[0061] Next, during step 106, a topic with the name “SEND_DATA” is created and a DataWriter (DW-0) is initialized during step 108 under P0 using the created topic. The reason for this hierarchy is that different topics (i.e. datasets) may be required to be sent independently by the same publisher, and each of these topics may have several DataWriters for redundancy. After that, the publisher during step 110 starts reading the seismic shots or simulation cells data from input, and initializes the data structure for the source sample (SS) by defining the matrices data size and the number of workers. The publisher then during step 110 starts sending the source sample SS-0 through the DataWriter DW-0. The procedure continues until step 112 indicates all sending has been completed. Processing then is continued at step 114.

[0062] As indicated at step 114, the second thread from master node 20 reverses the function of thread 0 by creating an instance of a subscriber S0 with selected QoS profile in Domain-0, in preparation to receive the partial results from the worker or processor nodes 24.

Specifically, a DataReader (DR-0) is configured during step 116 at master node 20 for the subscriber S0 (the workers) that uses topic "RECV_RESULT". Master node 20 then as indicated at step 118 listens to the workers through the receiving sample RS-0 and outputs the partial results. This processing continues as indicated at step 120 until receiving of sample RS-0 is completed.

[0063] On the subscribers' side (i.e., the workers), during step 122 each worker node 24 as discussed above and as shown in Figure 5 initiates itself as a subscriber to the main publisher P0, assigns an ID to itself (Wi), and starts receiving the seismic or simulation data for processing and subsequent result collection. The distribution of which data goes to which node is done dynamically in a way that is determined by first identifying the data size taken by each node according to the format of the data.

[0064] In case of a node failure of a processor node 24 on the workers' side, the system administrator of master node 20 may initiate a new processor node 24 with the same ID of the failed worker. The new worker would read the written checkpointed status as defined in the policy, re-read the sample from the persistence service 51, and resume the operation of the system.

[0065] It is important to mention that as a requirement for the durability QoS, all sent topics require DataWriters to match the configuration of the persistent QoS policy configuration with the DataReaders. As a consequence, a DataWriter that has an incompatible QoS with respect to what the topic specified will not send its data to the persistent service, and thus its status will not be saved. Similarly, a DataReader that has an incompatible QoS with respect to the specified in the topic will not get data from it.

[0066] Figures 6C and 6D indicate the methodology according to the present invention of providing either Deadline Quality of Service for real time applications checking of real time

data. (Figure 6C) or Liveliness Quality of Service (Figure 6D) for non-real time application checking of data. As mentioned, the selection is based on the determination made during step 152 (Figure 6A). During the data stream processing 100 shown in Figure 6B, either Deadline or Liveliness Quality of Service is continuously monitored according to the results of step 152.

[0067] If Deadline Quality of Service is so indicated, as shown at step 160 in Figure 6C, a determination is continuously made whether the Deadline QoS is met. If the deadline QoS is indicated not to be met during step 160, a specified time has passed during which real time data has not been published from publisher master node 20. If so, as indicated in Figure 6C, step 128 (Figure 7) is involved and a replacement master node 20 is substituted based on the highest Ownership QoS among the available master nodes 20, as will be described. If during the periodic Deadline QoS processing shown in Figure 6C real time data is being provided continuously by the publisher master node 20 according to the Deadline QoS and thus, within the specified period, the processing during step 100 of Figure 6C continues uninterrupted.

[0068] If during step 152, it is determined that the Liveliness QoS is to be in effect, a determination is continuously made during step 140 (Figures 6D and 7) whether the published master node is performing properly. If so, processing during step 100 continues uninterrupted as shown in Figure 6D. If not, a replacement master node 20 is substituted during step 128 based on the highest Ownership QoS among the available other master nodes 20, as will be described in connection with Figure 7.

[0069] Thus Figure 6 as well as Figures 6A through 6D illustrate an example parallel processing often encountered with exploration and production data according to the present invention. The results were as expected: QoS is controlled with fault-tolerance enabled when using the DDS techniques as adopted middleware in such computer processing of exploration

and production data. Specifically, compute nodes on the cluster are turned off and back on again, and the jobs continued to run with no need for a restart.

[0070] Figure 7 illustrates the role of the Ownership QoS, the Liveliness QoS, and the Deadline QoS in choosing one of the available master nodes 20 in the cluster. The master node 20 with the highest “OWNERSHIP_STRENGTH” number is chosen as the designated master publisher node 20 as indicated at step 128 based on an “ownership strength QoS” assigned as indicated at step 130. The assignment during step 130 is according to the criteria described.

[0071] Under control of the designated master publisher node 20, processing of the exploration and production data continues to process exploration and production data, as indicated as step 134 in the manner shown in Figure 5 until the processing job is completed, as indicated at step 136, so that the processed data can be sent to the publisher as shown at step 94 (Figure 5).

[0072] Periodically, based on the “Liveliness” quality of service profile, established in the manner set forth above, as indicated at step 138, an inquiry is made as indicated at step 140 of the operating status of the designated master publisher node 50. If it is determined during step 140 that the designated master publisher node 50 is not operating or functioning, a master node 20 having the next highest ownership strength number is introduced at the designated master publisher node according to step 128 and processing of the exploration and production data resumed in step 134. If during step 140, the present designated master publisher node 50 is indicated operable and functioning, processing according to step 134 continues.

[0073] The “LIVELINESS” QoS service according to the present invention thus continuously checks the health of the master node 50 in the absence of the Deadline QoS

service being in effect. This is done through sending status inquiry or heartbeat signals. If the node is unavailable, the “OWNERSHIP” QoS 130 is instructed to assign the next available master node with the highest “OWNERSHIP_STRENGTH” to continue the job. It is important to note that all available master nodes have access to the same data source. This is to allow them all to continue the stopped process.

[0074] As also shown in Figure 7, if the Deadline QoS service is invoked for real time processing, and such data is not being timely received, a replacement master node is substituted. As indicated at step 160 in Figures 6C and 7, if the deadline time specified for the Deadline is not met, step 128 is also invoked to assign the available master node 20 with the highest ownership strength number. Thereafter, step 134 is implemented in the manner described above.

[0075] The present invention provides the ability to control QoS properties on HPC and Grids that affect predictability, overhead, resource utilization, and aligns the scarce resources to the most critical requirements to these jobs.

[0076] To evaluate the performance of the present invention over conventional HPC and compare it with processing using MPI, data streaming was performed on the same data set using both paradigms (i.e., DDS with the present invention and MPI) and evaluated them on the data processing system clusters of Figures 1 and 2. The data streaming processing algorithm is computationally intensive with iterations, and it was chosen since it is a fundamental operation in many numerical linear algebra applications used in processing of exploration and production data. An efficient implementation on parallel computers is an issue of prime importance when providing such systems for processing of exploration and production data.

[0077] Figure 8 shows benchmarks to test the scalability and runtime of MPI and DDS by streamlining Reservoir Simulation data. It can be seen that the MPI version outperformed in terms of speed by taking around 6.76 seconds to stream 10GB size of file, compared with 7.9 seconds using DDS. This is expected since the QoS is adding more computations and checks to the communication level. As has been mentioned, however, MPI and PVM are focused on processing speed, with no effort to monitor accuracy or possible processing deficiencies.

[0078] Figure 9 shows the delay in engaging a new node according to the present invention, replacing a crashed node, while using the persistent service and durability, reliability and history QoS. This test is not applicable to a prior art MPI implementation. Since MPI implementations do not provide a capability of specifying service quality for computation or communication. As indicated in Figure 9, the delay in engaging a new node is proportional to the size of the matrices, since the persistent service needs to resend all the previously published instances to this new node. During the benchmark testing, test results depicted in Figure 8 indicate that it took 15.2 seconds in a 10GB data stream between nodes, while it took 72.2 seconds in a 50GB test. This is to be compared with the time required when it was necessary to resubmit the entire data set for processing for data in large quantities.

[0079] The present invention thus adds several benefits of having quality of service in high performance computing that are not available in the traditional method (i.e. by using MPI as a middleware). Among the benefits are that periodic publishers can indicate the speed at which they can publish by offering guaranteed update deadlines. By setting a deadline, a compliant publisher promises to send a new update at a minimum rate. Subscribers may then request data at that or any slower rate.

[0080] Another benefit is that the continuing participation or activity of entities can be monitored. The selected Liveliness and Durability QoS offered with the present invention determine whether an entity or a node is “active” (i.e., alive). The application can also be informed via a listener when an entity is no longer responsive. A further benefit is that the present invention also permits the data processing system D to automatically arbitrate between multiple publishers of the same topic with a parameter called "strength." Subscribers receive from the strongest active publisher. This provides automatic failover; if a strong publisher fails, all subscribers immediately receive updates from the backup (weaker) publisher.

[0081] It is also to be noted that QoS parameters exist with the DDS employment to control the resources of the entire system, suggest latency budgets, set delivery order, attach user data, prioritize messages, set resource utilization limits and partition the system into namespaces. The present invention thus provides the ability to control QoS properties on HPC and grids that affect predictability, overhead, resource utilization, and align the computational resources to the most critical requirements.

[0082] The present invention is a feasible option for those applications in which QoS is considered a priority, or for those HPC batch jobs that would run for several days on commodity hardware, where the probability of failure is not negligible. Accordingly, the present invention provides the ability to control QoS properties on HPC and grids that affect predictability, overhead, resource utilization, and aligns the scarce resources to the most critical requirements.

[0083] The invention has been sufficiently described so that a person with average knowledge in the matter may reproduce and obtain the results mentioned in the invention herein. Nonetheless, any skilled person in the field of technique, subject of the invention

herein, may carry out modifications not described in the request herein, to apply these modifications to a determined structure, or in the manufacturing process of the same, requires the claimed matter in the following claims; such structures shall be covered within the scope of the invention.

[0084] It should be noted and understood that there can be improvements and modifications made of the present invention described in detail above without departing from the spirit or scope of the invention as set forth in the accompanying claims.

CLAIMS

What is claimed is:

1. A computer implemented method of computerized processing in a data processing system of data for exploration and production of hydrocarbons, the data processing system including a plurality of master nodes, each with an established quality of service standard profile including an assigned ownership strength, with the master node of the plurality of master nodes having a highest assigned ownership strength being designated master publisher node for the exploration and production data being processed, the quality of service standard profile further including a specified deadline time interval for periodic production of data by the designated master publisher node when the data being processed is being produced in real time, the data processing system further including a plurality of processor nodes established as subscribers to receive exploration and production data from the designated master publisher node, and a data memory, the method comprising the computer processing steps of:

(a) transmitting the established quality of service standard profile from the designated master publisher node to the processor nodes;

(b) establishing with the designated master publisher node a domain for exploration and production data processing by the designated master publisher node and designated ones of the plurality of processor nodes as subscriber processor nodes;

(c) producing the exploration and production data from the designated master publisher node to the designated subscriber processor nodes of the domain;

(d) determining if the data being processed is indicated to be real time data;

(e) processing the received exploration and production data in the designated subscriber processor nodes of the domain;

(f) monitoring at the designated master publisher node the processed exploration and production data of the designated subscriber processor nodes of the domain;

(g) assembling in the data memory of the data processing system the processed exploration and production data received at the designated master publisher node;

(h) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine if the specified deadline time interval has passed without periodic production of data from the designated master publisher node; and

(i) if the specified deadline time interval has not passed, continuing the step of monitoring performance of the designated master publisher node to determine that the designated master publisher node is operating, continuing the processing of the exploration and production data; and

(j) if the designated master publisher node is not operating, determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(k) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

2. The computer implemented method of Claim 1, wherein the established quality of service standard profile further includes a liveliness quality of service control to periodically determine the operating status of the designated master publisher node, and wherein the step of determining if the data being processed is indicated to be real time data, and further including the steps of:

(a) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine operating status of the designated master publisher node;

(b) if the step of monitoring performance of the designated master publisher node indicates the designated master publisher node is operating, continuing the processing of the exploration and production data; and, if not

(c) determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(d) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

3. The computer implemented method of Claim 1, wherein the data processing system further includes a data display, and further including the computer processing step of:

forming an output display of the assembled processed exploration and production data.

4. The computer implemented method of Claim 1, wherein the exploration and production data comprises a reservoir simulation model.
5. The computer implemented method of Claim 1, wherein the exploration and production data comprises a geological model of a reservoir.
6. The computer implemented method of Claim 1, wherein the exploration and production data comprises a seismic survey of the earth in a reservoir.
7. The computer implemented method of Claim 1, further including the steps of:

(a) determining in the designated master publisher node whether the designated subscriber processor nodes of the domain comply with the transmitted established quality of service standard profile from the designated master publisher node; and

(b) if so, receiving at the designated master publisher node the processed exploration and production data from the designated subscriber processor nodes which comply with the transmitted established quality of service standard profile; and

(c) if not, inhibiting at the designated master publisher node transfer to the designated master publisher of the processed exploration and production data from the designated subscriber processor nodes which do not comply with the transmitted established quality of service standard profile;

(d) joining additional processor nodes as designated subscriber processor nodes of the domain during processing of exploration and production data; and

(e) delivering the produced exploration and production data to the joined additional processor nodes.

8. A data processing system for computerized processing of data for exploration and production of hydrocarbons, the data processing system comprising a plurality of master nodes, each with an established quality of service standard profile including an assigned ownership strength, with the master node of the plurality of master nodes having a highest assigned ownership strength being designated master publisher node for the exploration and production data being processed, the quality of service standard profile further including a specified deadline time interval for periodic production of data by the designated master publisher node when the data being processed is being produced in real time, a plurality of processor nodes and a data memory, the data processing system further comprising:

(a) the designated master publisher node performing the steps of:

(1) transmitting the established quality of service standard profile from the designated master publisher node to the processor nodes;

(2) establishing with the designated master publisher node a domain for exploration and production data processing by the designated master publisher node and designated ones of the plurality of processor nodes as subscriber processor nodes;

(3) producing the exploration and production data from the designated master publisher node to the designated subscriber processor nodes of the domain;

(4) determining if the data being processed is indicated to be real time data;

(b) the plurality of processor nodes established as subscribers to receive exploration and production data from the designated master publisher node, the plurality of processor nodes performing the steps of:

(1) receiving in the designated subscriber processor nodes of the domain the exploration and production data and the established quality of service standard profile from the designated master publisher node;

(2) processing the received exploration and production data in the designated subscriber processor nodes of the domain; and

(c) the designated master publisher node further performing the steps of:

(1) monitoring the processed exploration and production data of the designated subscriber processor nodes of the domain;

(2) assembling in the data memory the processed exploration and production data from the designated subscriber processor nodes which comply with the transmitted established quality of service standard profile;

(3) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine if the specified deadline time interval has passed without periodic production of data from the designated master publisher node;

(4) if the specified deadline time interval has not passed, continuing the step of monitoring performance of the designated master publisher node to determine that the designated master publisher node is operating, continuing the processing of the exploration and production data;

(5) if the designated master publisher node is not operating, determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(6) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

9. The data processing system of Claim 8, wherein the established quality of service standard profile further includes a liveness quality of service control to periodically determine the operating status of the designated master publisher node, and wherein the step of determining if the data being processed is indicated to be real time data, and wherein the processor nodes further performs the steps of:

(a) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine operating status of the designated master publisher node;

(b) if the step of monitoring performance of the designated master publisher node indicates the designated master publisher node is operating, continuing the processing of the exploration and production data; and, if not

(c) determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(d) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

10. The data processing system of Claim 8 further including a data display, and wherein the master node further performs the step of:

forming an output display of the assembled processed exploration and production data.

11. The data processing system of Claim 8, wherein the exploration and production data comprises a reservoir simulation model.

12. The data processing system of Claim 8, wherein the exploration and production data comprises a geological model of a reservoir.

13. The data processing system of Claim 8, wherein the exploration and production data comprises a seismic survey of the earth in a reservoir.

14. The data processing system of Claim 8, wherein the processor nodes further perform the steps of:

determining whether the designated subscriber processor nodes of the domain comply with the transmitted established quality of service standard profile from the designated master publisher node;

if so, receiving the processed exploration and production data from the designated subscriber processor nodes which comply with the transmitted established quality of service standard profile;

if not, inhibiting at the designated master publisher node transfer of the processed exploration and production data from the designated subscriber processor nodes which do not comply with the transmitted established quality of service standard profile;

joining additional processor nodes as designated subscriber processor nodes of the domain during processing of exploration and production data;

delivering the produced exploration and production data to the joined additional processor nodes.

15. A data storage device having stored in a non-transitory computer readable storage medium computer operable instructions for causing a data processing system to perform computerized processing of data for exploration and production of hydrocarbons, the data processing system including a plurality of master nodes, each with an established quality of service standard profile including an assigned ownership strength, and with the master node of the plurality of master nodes having a highest assigned ownership strength being designated master publisher node for the exploration and production data being processed, the quality of service standard profile further including a specified deadline time interval for

periodic production of data by the designated master publisher node when the data being processed is being produced in real time, the data processing system including a plurality of processor nodes established as subscribers to receive the exploration and production data from the designated master publisher node, and a data memory, the instructions stored in the data storage device causing the data processing system to perform the following steps:

- (a) transmitting the established quality of service standard profile from the designated master publisher node to the processor nodes;
- (b) establishing with the designated master publisher node a domain for exploration and production data processing by the designated master publisher node and designated ones of the plurality of processor nodes as subscriber processor nodes;
- (c) producing the exploration and production data from the designated master publisher node to the designated subscriber processor nodes of the domain;
- (d) determining if the data being processed is indicated to be real time data;
- (e) processing the received exploration and production data in the designated subscriber processor nodes of the domain;
- (f) monitoring at the designated master publisher node the processed exploration and production data of the designated subscriber processor nodes of the domain;
- (g) assembling in the data memory of the data processing system the processed exploration and production data received at the designated master publisher node;
- (h) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine if the specified deadline time

interval has passed without periodic production of data from the designated master publisher node; and

(i) if the specified time deadline interval has not passed, continuing the step of monitoring performance of the designated master publisher node to determine that the designated master publisher node is operating, continuing the processing of the exploration and production data; and,

(j) if the designated master publisher node is not operating, determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(k) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

16. The data storage device of Claim 15, wherein the established quality of service standard profile further includes a liveliness quality of service control to periodically determine the operating status of the designated master publisher node, and wherein the data being processed is indicated to be real time data, and further include instructions to perform the steps of:

(a) monitoring performance of the designated master publisher node during the processing of the exploration and production data to determine operating status of the designated master publisher node;

(b) if the step of monitoring performance of the designated master publisher node indicates the designated master publisher node is operating, continuing the processing of the exploration and production data; and, if not

(c) determining ownership strength of the others of the plurality of master nodes according to their respective ownership strength quality of service standard profile; and

(d) establishing as designated master publisher node the master node having the highest ownership strength quality of service standard profile and resuming the processing of the exploration and production data.

17. The data storage device of Claim 15, wherein the data processing system further includes a data display, and further including instructions for performing the step of:

forming an output display of the assembled processed exploration and production data.

18. The data storage device of Claim 15, wherein the exploration and production data comprises a reservoir simulation model.

19. The data storage device of Claim 15, wherein the exploration and production data comprises a geological model of a reservoir.

20. The data storage device of Claim 15, wherein the exploration and production data comprises a seismic survey of the earth in a reservoir.

21. The data storage device of Claim 15, further including instructions for performing the steps of:

(a) determining in the designated master publisher node whether the designated subscriber processor nodes of the domain comply with the transmitted established quality of service standard profile from the designated master publisher node; and

(b) if so, receiving at the designated master publisher node the processed exploration and production data from the designated subscriber processor nodes which comply with the transmitted established quality of service standard profile; and

(c) if not, inhibiting at the designated master publisher node transfer to the designated master publisher node of the processed exploration and production data from the designated subscriber processor nodes which do not comply with the transmitted established quality of service standard profile;

(d) joining additional processor nodes as designated subscriber processor nodes of the domain during processing of exploration and production data; and

(e) delivering the produced exploration and production data to the joined additional processor nodes.




FIG. 1
(PRIOR ART)




FIG. 2

2/8

FIG. 3





FIG. 5



3/8




FIG. 4

4/8

FIG. 6






FIG. 6A



5/8

FIG. 6B



6/8




FIG. 6C




FIG. 6D




FIG. 7

8/8

FIG. 8





FIG. 9



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/039760

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F9/50 G06F11/14
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/362619 A1 (AISHAIKH RAED ABDULLAH [SA] ET AL) 17 December 2015 (2015-12-17) abstract paragraphs [0004], [0006], [0016], [0039], [0050], [0071], [0072]; claims 13,27; figures 2,7 -----	1-21



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 September 2017

Date of mailing of the international search report

26/09/2017

Name and mailing address of the ISA/

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

Authorized officer

Kingma, Ype

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/039760

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015362619	A1	17-12-2015	NONE
