

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 August 2011 (18.08.2011)

PCT

(10) International Publication Number
WO 2011/100002 A1

(51) International Patent Classification:
G06G 7/48 (2006.01)

(21) International Application Number:
PCT/US2010/053141

(22) International Filing Date:
19 October 2010 (19.10.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/304,056 12 February 2010 (12.02.2010) US

(71) Applicant (for all designated States except US):
EXXONMOBIL UPSTREAM RESEARCH COMPANY [US/US]; P.O. Box 2189 (CORP-URC-SW359), Houston, TX 77252-2189 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **MALIASSOV, Sergei** [US/US]; 21003 La Arbore Lane, Spring, TX 77388 (US). **SHUTTLEWORTH, Robert, R.** [US/US]; 624 Westcott Street, Houston, TX 77007 (US).

(74) Agents: **JENSEN, Nathan, O.** et al.; ExxonMobil Upstream Research Company, P.O. Box 2189 (CORP-URC-SW359), Houston, TX 77252-2189 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR PARTITIONING PARALLEL SIMULATION MODELS



100
FIG. 1

(57) Abstract: A method is presented for partitioning a simulation model into a plurality of subdomains that may each be assigned to one of a plurality of processors. The method includes creating a representation of a topology graph of a simulation model in a tangible, computer readable medium. The topology graph includes a plurality of computational elements and a plurality of connections between those elements. Each of the plurality of connections is weighted to create a plurality of weights, and each of the plurality of weights is scaled. Optionally, the weights can be mapped to different interval of values. Based on the weights information the topology graph is partitioned into two or more subdomains, wherein a partition boundary follows a local topographical minimum in the topology graph. A subdomain is assigned to each of the plurality of processors.



WO 2011/100002 A1

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
 - *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*
- Published:**
- *with international search report (Art. 21(3))*

Attorney No.: 2010EM054

**METHOD AND SYSTEM FOR
PARTITIONING PARALLEL SIMULATION MODELS**

CROSS-REFERENCE TO RELATED APPLICATION

5

[0001] This application claims the benefit of U.S. Provisional Application No. 61/304,056 filed February 12, 2010, entitled Method and System for Partitioning Parallel Simulation Models, which is incorporated by reference, in its entirety, for all purposes.

FIELD

10 [0002] Exemplary embodiments of the present techniques relate to a method and system for partitioning parallel simulation models between computer systems.

BACKGROUND

15 [0003] This section is intended to introduce various aspects of the art, which may be associated with exemplary embodiments of the present techniques. This discussion is believed to assist in providing a framework to facilitate a better understanding of particular aspects of the present techniques. Accordingly, it should be understood that this section should be read in this light, and not necessarily as admissions of prior art.

20 [0004] Modern society is greatly dependant on the use of hydrocarbons for fuels and chemical feedstocks. Hydrocarbons are generally found in subsurface rock formations that can be termed "reservoirs." Removing hydrocarbons from the reservoirs depends on numerous physical properties of the rock formations, such as the permeability of the rock containing the hydrocarbons, the ability of the hydrocarbons to flow through the rock formations, and the proportion of hydrocarbons present, among others.

25 [0005] Often, mathematical models termed "simulation models" are used to simulate hydrocarbon reservoirs and optimize the production of the hydrocarbons. A simulation model is a type of computational fluid dynamics simulation where a set of partial differential equations (PDE's) which govern multi-phase, multi-component fluid flow through porous media and the connected facility network is approximated and solved. This is an iterative, time-stepping process where a particular hydrocarbon production strategy is optimized.

30 [0006] Simulation models discretize the underlying PDEs on a structured (or unstructured) grid, which represents the reservoir rock, wells, and surface facility network.

Attorney No.: 2010EM054

State variables, such as pressure and saturation, are defined at each grid block. The goal of a simulation model is generally to understand the flow patterns of the underlying geology in order to optimize the production of hydrocarbons from a set of wells and surface facilities. During the past five decades, the size and complexity of simulation models have grown
5 proportionally with the increased availability of computing capacity. Complex simulation models often require the use of parallel computing systems and algorithms to provide adequate simulation turnaround time.

[0007] Accordingly, a simulation model can be divided into a number of subsections, or subdomains, wherein each subdomain may be assigned or partitioned to a different
10 computing unit, such as a processor in a cluster computing system or a computing core in a multi-core desktop. However, poor partitioning may result in slow convergence to a solution or even in a failure of the model to converge to a solution. For example, if a partition of a simulation model between different computing units crosses a production well, the simulation may fail. Generally, partitioning may be performed by a mathematical analysis of a graph
15 representing the simulation model.

[0008] Recently, there have been a number of research articles published on the general subject of graph partitioning, for example, as applied to finite element analysis and other problems. See, for example, J.D. Teresco, K.D. Devine, and J.E. Flaherty, "Partitioning and Dynamic Load Balancing for the Numerical Solution of Partial Differential Equations,"
20 Numerical Solution of Partial Differential Equations on Parallel Computers, Chapter 2, 55–88 (Springer, New York, 2006); and B. Hendrickson and T. Kolda, "Graph Partitioning Models for Parallel Computing," Parallel Computing, 26, 1519-1534 (2000). Some of the partitioning techniques have been implemented in software, such as the METIS program. See G. Karypis and V. Kumar, "A fast and high quality multilevel scheme for partitioning
25 irregular graphs," SIAM Journal on Scientific Computing, 20, 359-392 (1999). Another example is the Chaco program, developed at Sandia National Laboratory. See B. Hendrickson and R. Leland, "An Improved Spectral Graph Partitioning Algorithm for Mapping Parallel Computations," SIAM Journal on Scientific and Statistical Computing, 16, 452-469 (1995); see also B. Hendrickson and R. Leland, "The Chaco User's Guide: Version
30 2.0," Sandia Tech Report SAND94-2692 (1994).

Attorney No.: 2010EM054

[0009] The graph partitioning algorithms are organized into two main classes: global and multi-level. The global algorithms take as input a global graph, or map of the simulation, and perform a spectral technique (such as an eigenvalue decomposition of a Laplacian representation of the global graph) to generate the partitions. These algorithms produce good quality partitions, but are computationally expensive for large problems. Both Chaco and METIS have global algorithms implemented in each software package.

[0010] In contrast, multi-level algorithms take the global graph as an input, but create a sequence of coarser, or less detailed, graphs until a certain threshold is reached. The coarsest graph is partitioned using a local optimization technique, such as a spectral technique. The partition on the coarsest level is projected back to finer levels, thus creating a partition for the entire grid. For large graphs, multi-level techniques are preferable due to faster computational time than the global techniques.

[0011] Software packages that implement partitioning algorithms generally allow the user to assign a weight to the connection associated with each graph node. For most algorithms, the weights are positive integers. The integer values of the weights should have a relatively short range, such as from about 1 to about 1000 or about 10,000. As the range increases, the robustness of a partitioning algorithm and the quality of a created partition deteriorates. Moreover, the published algorithms and software do not include methods for keeping certain grid blocks or physical properties of a model together in one subdomain. In other words, when using publicly available software, such as METIS or Chaco it is impossible to ensure that certain graph connections are preserved resulting in the corresponding nodes being partitioned into separate subdomains by the partitioning algorithm.

[0012] There are approaches to assign the real weight values to edges and nodes of the graph to improve the quality of the partition. Some of them, pertinent to problems in the oil and gas industry are described in International Patent Publication No. WO2009/075945. The patent describes how to construct real weights for nodes and/or connections based on physical information (for example, transmissibility, flux values, etc). Unfortunately, there is not a robust algorithm to map real weights with a very large range to a short range of integer values, in order to produce a robust partition for simulation models.

30

Attorney No.: 2010EM054

SUMMARY

[0013] An exemplary embodiment of the present techniques provides a method for partitioning a reservoir, for example, to distribute the workload of a simulation model among processors in a computing cluster. The method may include generating a representation of a topology graph of a simulation model in a tangible, computer readable medium, wherein the topology graph comprises a plurality of elements and a plurality of connections between adjacent elements. Each of the plurality of connections may be weighted to create a plurality of weights and each of the plurality of weights may be scaled. The topology graph may be partitioned into two or more subdomains, wherein a partition boundary follows a local topographical minimum in the topology graph. A subdomain may be assigned to each of a plurality of processors.

[0014] In some embodiments, the plurality of elements may include computational cells in a computational mesh. The plurality of elements may include rows in a linear system matrix and the plurality of connections correspond to nonzero elements of the matrix.

[0015] In some embodiments, the weighting of each of the plurality of connections may be based, at least in part, on physical properties assigned to the computational cells. The physical properties may include transmissibility, total mobility, mass flow, heat flow, or any combinations thereof. The weighting of each of the plurality of connections may be based at least in part on a plurality of off-diagonal coefficients in a Jacobian matrix representing a simulation model. The weighting of each of the plurality of connections may be based, at least in part, upon its proximity to a well bore. A near-well region may be kept in one of the two or more subdomains. The weighting of each of the plurality of connections may be based, at least in part, upon its belonging to a grid block that is perforated by a well bore. Further, the weighting of each of the plurality of connections may be based, at least in part, upon a solution to a local flow problem.

[0016] In some embodiments, scaling each of the plurality of weights may be performed using a probability distribution. Further, scaling each of the plurality of weights may be performed by scaling discrete reservoir properties to create a linear mapping. In exemplary embodiments, the simulation model may be performed for a subdomain assigned to each of the plurality of processors.

Attorney No.: 2010EM054

[0017] Another exemplary embodiment provides a system for modeling reservoir properties. The system may include a plurality of processors and a storage medium comprising a representation of a topology graph of a reservoir model, wherein the topology graph comprises a plurality of connections. The system may also include a machine readable
5 medium comprising code configured to direct at least one of the plurality of processors to weight each of the plurality of connections to create a plurality of weights, map each of the plurality of weights to an integer value, partition the topology graph into two or more subdomains, and assign each of the two or more subdomains to one of the plurality of processors.

10 [0018] In some embodiments, each of the plurality of weights may be based, at least in part, on an associated physical property. The physical property may represent heat transfer, mass transfer, total flow, transmissibility, or any combinations thereof. The plurality of processors may include a cluster computing system.

[0019] Another exemplary embodiment provides a tangible, computer readable medium
15 that includes code configured to direct a processor to create a topology graph of a simulation model, wherein the topology graph comprises a plurality of connections between the center points of adjacent computational cells in a computational mesh. The code may also be configured to weight each of the plurality of connections to create a plurality of weights and scale each of the plurality of weights. The code may further be configured to partition the
20 topology graph into two or more subdomains, wherein a partition boundary follows a local topographical minimum in the topology graph. The code may be configured to assign each of the subdomains to one of a plurality of processors.

[0020] In some embodiments, the code may be configured to direct the one of the plurality of processors to process a simulation model for the subdomain. Further, the code
25 may be configured to direct the processor to map the plurality of weights into an integer range.

DESCRIPTION OF THE DRAWINGS

[0021] The advantages of the present techniques are better understood by referring to the following detailed description and the attached drawings, in which:

30 [0022] Fig. 1 is a schematic view of a reservoir, in accordance with an exemplary embodiment of the present techniques;

Attorney No.: 2010EM054

[0023] Fig. 2 is a top view of a reservoir showing a planar projection of a computational mesh over the reservoir, in accordance with an exemplary embodiment of the present techniques;

5 [0024] Fig. 3 is a close up view of the computational mesh, illustrating connections between computational cells, in accordance with an exemplary embodiment of the present techniques;

[0025] Fig. 4 is a process flow diagram of a workflow for modelling a reservoir, in accordance with an exemplary embodiment of the present techniques;

10 [0026] Fig. 5 is a block diagram of a method for partitioning a simulation model, in accordance with an exemplary embodiment of the current techniques;

[0027] Fig. 6 is a block diagram of a method for partitioning a simulation model, in accordance with an exemplary embodiment of the present techniques;

[0028] Fig. 7 is a process flow diagram illustrating a method for scaling connection weights, in accordance with an exemplary embodiment of the present techniques;

15 [0029] Fig. 8 is a graph illustrating a mapping of real values to integer weights using a cumulative distribution function, in accordance with an exemplary embodiment of the present techniques; and

[0030] Fig. 9 is a block diagram of an exemplary cluster computing system that may be used in exemplary embodiments of the present techniques.

20

DETAILED DESCRIPTION

[0031] In the following detailed description section, the specific embodiments of the present techniques are described in connection with preferred embodiments. However, to the extent that the following description is specific to a particular embodiment or a particular use of the present techniques, this is intended to be for exemplary purposes only and simply provides a description of the exemplary embodiments. Accordingly, the present techniques are not limited to the specific embodiments described below, but rather, such techniques include all alternatives, modifications, and equivalents falling within the true spirit and scope of the appended claims.

25

Attorney No.: 2010EM054

[0032] At the outset, and for ease of reference, certain terms used in this application and their meanings as used in this context are set forth. To the extent a term used herein is not defined below, it should be given the broadest definition persons in the pertinent art have given that term as reflected in at least one printed publication or issued patent. Further, the present techniques are not limited by the usage of the terms shown below, as all equivalents, synonyms, new developments, and terms or techniques that serve the same or a similar purpose are considered to be within the scope of the present claims.

[0033] “Coarsening” refers to reducing the number of cells in simulation models by making the cells larger, for example, representing a larger space in a reservoir. Coarsening is often used to lower the computational costs by decreasing the number of cells in a geologic model prior to generating or running simulation models.

[0034] “Computer-readable medium” or “tangible, computer-readable medium” as used herein refers to any tangible storage and/or transmission medium that participates in providing instructions to a processor for execution. Such a medium may include, but is not limited to, non-volatile media and volatile media. Non-volatile media includes, for example, NVRAM, or magnetic or optical disks. Volatile media includes dynamic memory, such as main memory. Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, a hard disk, an array of hard disks, a magnetic tape, or any other magnetic medium, magneto-optical medium, a CD-ROM, a holographic medium, any other optical medium, a RAM, a PROM, and EPROM, a FLASH-EPROM, a solid state medium like a memory card, any other memory chip or cartridge, or any other tangible medium from which a computer can read data or instructions. When the computer-readable media is configured as a database, it is to be understood that the database may be any type of database, such as relational, hierarchical, object-oriented, and/or the like.

[0035] As used herein, “to display” or “displaying” includes a direct act that causes displaying, as well as any indirect act that facilitates displaying. Indirect acts include providing software to an end user, maintaining a website through which a user is enabled to affect a display, hyperlinking to such a website, or cooperating or partnering with an entity who performs such direct or indirect acts. Thus, a first party may operate alone or in cooperation with a third party vendor to enable the reference signal to be generated on a display device. The display device may include any device suitable for displaying the

Attorney No.: 2010EM054

reference image, such as without limitation a CRT monitor, a LCD monitor, a plasma device, a flat panel device, or printer. The display device may include a device which has been calibrated through the use of any conventional software intended to be used in evaluating, correcting, and/or improving display results (for example, a color monitor that has been adjusted using monitor calibration software). Rather than (or in addition to) displaying the reference image on a display device, a method, consistent with the invention, may include providing a reference image to a subject. "Providing a reference image" may include creating or distributing the reference image to the subject by physical, telephonic, or electronic delivery, providing access over a network to the reference, or creating or distributing software to the subject configured to run on the subject's workstation or computer including the reference image. In one example, the providing of the reference image could involve enabling the subject to obtain the reference image in hard copy form via a printer. For example, information, software, and/or instructions could be transmitted (for example, electronically or physically via a data storage device or hard copy) and/or otherwise made available (for example, via a network) in order to facilitate the subject using a printer to print a hard copy form of reference image. In such an example, the printer may be a printer which has been calibrated through the use of any conventional software intended to be used in evaluating, correcting, and/or improving printing results (for example, a color printer that has been adjusted using color correction software).

20 [0036] "Exemplary" is used exclusively herein to mean "serving as an example, instance, or illustration." Any embodiment described herein as "exemplary" is not to be construed as preferred or advantageous over other embodiments.

25 [0037] "Flow simulation" is defined as a numerical method of simulating the transport of mass (typically fluids, such as oil, water and gas), energy, and momentum through a physical system using a computer. The physical system includes a three dimensional reservoir model, fluid properties, the number and locations of wells. Flow simulations also require a strategy (often called a well-management strategy) for controlling injection and production rates. These strategies are typically used to maintain reservoir pressure by replacing produced fluids with injected fluids (for example, water and/or gas). When a flow simulation correctly recreates a past reservoir performance, it is said to be "history matched," and a higher degree of confidence is placed in its ability to predict the future fluid behavior in the reservoir.

30

Attorney No.: 2010EM054

[0038] “Permeability” is the capacity of a rock to transmit fluids through the interconnected pore spaces of the rock. Permeability may be measured using Darcy's Law: $Q = (k \Delta P A) / (\mu L)$, wherein Q = flow rate (cm^3/s), ΔP = pressure drop (atm) across a cylinder having a length L (cm) and a cross-sectional area A (cm^2), μ = fluid viscosity (cp), and k = permeability (Darcy). The customary unit of measurement for permeability is the millidarcy. The term “relatively permeable” is defined, with respect to formations or portions thereof, as an average permeability of 10 millidarcy or more (for example, 10 or 100 millidarcy). The term “relatively low permeability” is defined, with respect to formations or portions thereof, as an average permeability of less than about 10 millidarcy. An impermeable layer generally has a permeability of less than about 0.1 millidarcy.

[0039] “Monotone mapping” is a function or relation between two sets of real numbers that preserves order. This mapping is a relation between an input real number (or integer) to an output real number (or integer) so that the relation is consistently increasing or decreasing.

[0040] “Pore volume” or “porosity” is defined as the ratio of the volume of pore space to the total bulk volume of the material expressed in percent. Porosity is a measure of the reservoir rock's storage capacity for fluids. Porosity is preferably determined from cores, sonic logs, density logs, neutron logs or resistivity logs. Total or absolute porosity includes all the pore spaces, whereas effective porosity includes only the interconnected pores and corresponds to the pore volume available for depletion.

[0041] “Reservoir” or “reservoir formations” are typically pay zones (for example, hydrocarbon producing zones) that include sandstone, limestone, chalk, coal and some types of shale. Pay zones can vary in thickness from less than one foot (0.3048 m) to hundreds of feet (hundreds of m). The permeability of the reservoir formation provides the potential for production.

[0042] “Reservoir properties” and “reservoir property values” are defined as quantities representing physical attributes of rocks containing reservoir fluids. The term “reservoir properties” as used in this application includes both measurable and descriptive attributes. Examples of measurable reservoir property values include porosity, permeability, water saturation, and fracture density. Examples of descriptive reservoir property values include facies, lithology (for example, sandstone or carbonate), and environment-of-deposition

Attorney No.: 2010EM054

(EOD). Reservoir properties may be populated into a reservoir framework to generate a reservoir model.

[0043] “Simulation model” refers to a specific mathematical representation of a real hydrocarbon reservoir, which may be considered to be a particular type of geologic model.

5 Simulation models are used to conduct numerical experiments (reservoir simulations) regarding future performance of the field with the goal of determining the most profitable operating strategy. An engineer managing a hydrocarbon reservoir may create many different simulation models, possibly with varying degrees of complexity, in order to quantify the past performance of the reservoir and predict its future performance.

10 [0044] “Transmissibility” refers to the volumetric flow rate between two points at unit viscosity for a given pressure-drop. Transmissibility is a useful measure of connectivity. Transmissibility between any two compartments in a reservoir (fault blocks or geologic zones), or between the well and the reservoir (or particular geologic zones), or between injectors and producers, can all be useful for understanding connectivity in the reservoir.

15 [0045] “Well” or “wellbore” includes cased, cased and cemented, or open-hole wellbores, and may be any type of well, including, but not limited to, a producing well, an experimental well, an exploratory well, and the like. Wellbores may be vertical, horizontal, any angle between vertical and horizontal, deviated or non-deviated, and combinations thereof, for example a vertical well with a non-vertical component. Wellbores are typically drilled and
20 then completed by positioning a casing string within the wellbore. Conventionally, the casing string is cemented to the well face by circulating cement into the annulus defined between the outer surface of the casing string and the wellbore face. The casing string, once embedded in cement within the well, is then perforated to allow fluid communication between the inside and outside of the tubulars across intervals of interest. The perforations
25 allow for the flow of treating chemicals (or substances) from the inside of the casing string into the surrounding formations in order to stimulate the production or injection of fluids. Later, the perforations are used to receive the flow of hydrocarbons from the formations so that they may be delivered through the casing string to the surface, or to allow the continued injection of fluids for reservoir management or disposal purposes.

30

Attorney No.: 2010EM054

Overview

[0046] Exemplary embodiments of the present techniques disclose methods and systems for partitioning simulation models among computing units, such as individual computer systems, computing units in a computer cluster, and the like. The techniques are designed to minimize the chance of placing a partition separating two computational subdomains through a near well region or geologic feature (such as faults). In an exemplary embodiment, the techniques map real valued connection weights onto a narrow range of integer weights for use in partitioning algorithms, such as the publically available METIS or Chaco algorithms.

[0047] Fig. 1 is a schematic view 100 of a reservoir 102, in accordance with an exemplary embodiment of the present techniques. The reservoir 102, such as an oil or natural gas reservoir, can be a subsurface formation that may be accessed by drilling wells 104, 106, and 108 from the surface 110 through layers of overburden 112. The reservoir 102 may have one or more faults 114 dividing areas, for example regions 116 and 118, and which may either restrict or enhance the flow of hydrocarbons. The wells 104, 106, and 108 may be deviated, such as being directionally drilled to follow the reservoir 102. Further, the wells can be branched to increase the amount of hydrocarbon that may be drained from the reservoir, as shown for wells 104 and 108. The wells 104, 106, and 108, can have numerous areas with perforations 120 (indicated as dots next to the wells) to allow hydrocarbons to flow from the reservoir 102 into the wells 104, 106, and 108 for removal to the surface.

[0048] A simulation model, or simulator, of the reservoir 102 is likely to find that the greatest changes occur in the vicinity of the wells 104, 106, and 108, and other reservoir features, such as the fault 114. Accordingly, it would be useful to keep areas in the vicinity of each of these features in single computational subdomains. A partition between computational subdomains that crosses a well 104, 106, and 108, fault 114, or other feature may slow convergence of the simulation, increase computational loading by increasing communication between computing units, or even prevent convergence, resulting in a failure to find a solution.

[0049] Fig. 2 is a top view of a reservoir showing a planar projection of a computational mesh 200 over the reservoir, in accordance with an exemplary embodiment of the present techniques. Although the computational mesh 200 is shown as a two dimensional grid of computational cells (or blocks) 202 to simplify the explanation of the problem, it should be

Attorney No.: 2010EM054

understood that the actual computational mesh 200 can be a three dimension matrix of computational cells 202 that encompasses the reservoir. A computational cell 202 is a single two or three dimensional location within a simulation model that represents a physical location in a reservoir. The computational cell 202 may have associated properties, such as a porosity or an oil content, which is assumed to be a single value over the entire computational cell 202 and is assigned to the center of the computational cell 202. Computational cells 202 may interact with adjacent computational cells 202, for example, by having flux properties assigned to a shared border with the adjacent computational cells 202. For example, the flux properties may include heat or mass transfer values.

5
10 **[0050]** The computational mesh 200 can be coarsened in areas that may have less significant changes, for example, by combining computational cells 202 that are not in proximity to a well or other reservoir feature. Similarly, the computational mesh 200 may retain a fine mesh structure in the vicinity of wells or other reservoir features, such as the first well 204, or other reservoir features, for example, a second well 206, a third well 208, a fault
15 210, or any other features that may show larger changes than other areas.

[0051] The computational mesh 200 represents the simulation model, and can be divided among computing units to decrease the amount of time needed to provide a result for the simulation. This procedure may be termed “parallelization.” The parallelization of the simulation model is realized by parallelizing the individual components at each time step. To
20 achieve efficient utilization of the parallel computing units the simulation model can be distributed across the computing units so that the computational load is evenly balanced and the amount of inter-unit communication is minimized. This division is performed by partitioning the simulation model, i.e., assigning different computational cells 202 in the computational mesh 200 to different computing units (such as described with respect to Fig.
25 9). Each computational cell 202 may require a different approach to parallelization based on the numerical formulation, the actual input data, the computational task, and user supplied options.

[0052] In the exemplary embodiment shown in Fig. 2, the computational mesh 200 is partitioned between four computing units, as indicated by the subdomains labeled I-IV.
30 Although four computing units are used in Fig. 2, any number of computing units may be used in other embodiments, depending on the size of the simulation model and the number of

Attorney No.: 2010EM054

near well features. For example, a small simulation model may provide results in a reasonable timeframe from a single computing device, while a large simulation may use 10, 100, 1000, or even more computing units for the parallelization.

5 [0053] Further, while the subdomains I-IV do not cross near well regions or significant reservoir features, the subdomains are not limited to contiguous areas, but may include non-contiguous areas, which may be useful for balancing the load between the computing units. For example, as illustrated in Fig. 2, subdomain I may be divided into two regions. A first region 212 encompasses the near well region for the first well 204, while a second region 214 encompasses a number of larger computational cells 202 that may have less significant
10 changes than the near well regions.

[0054] Fig. 3 is a close up view of the computational mesh 200, illustrating connections between computational cells 202, in accordance with an exemplary embodiment of the present techniques. As shown in Fig. 3, each computational cell 202, has a center point 302, which may have associated static properties, such as oil content, permeability, and pressure,
15 among others. A connection 304 can be defined between the center point 302 of each computational cell 202. The connections 302 may also have associated flux properties, such as mass flow or heat flow.

[0055] The connections 304 form a topology graph of the grid connectivity which may be used for partitioning the computational mesh 200. Further, the weighting of the connections
20 304 may be used to determine whether a partitioning algorithm will be allowed to cut through the connection 304 and send computational cells 202 on one side to a first computing unit and computational cells on the other side to a second computing unit.

[0056] The simulation model and the partitioning procedure may be more clearly understood by examining a simulation workflow, as discussed with respect to Fig. 4.
25 Exemplary embodiments of the present techniques showing the partitioning of a simulation model are discussed further with respect to Figs. 5 and 6. A method for weighting a connection 304 that may be used in exemplary embodiments is discussed with respect to Figs. 7 and 8. A computing device that may be used in exemplary embodiments of the present techniques is discussed with respect to Fig. 9.

Attorney No.: 2010EM054

Workflow for Modelling a Reservoir

[0057] Fig. 4 is a process flow diagram of a workflow 400 for modelling a reservoir, in accordance with an exemplary embodiment of the present techniques. Although the discretization (coarsening) and the level of implicitness (which state variables, such as pressure or saturation, are treated implicitly or explicitly in the formulation) of the solution process varies, simulation models may perform in a similar fashion as workflow 400. A simulation model can begin at block 402 by parsing user input data. The input data may include the problem formulation, a geologic model that is discretized into grid blocks with physical properties defined at each grid block, including rock properties (such as permeability) and fluid properties (such as transmissibility). At block 404, a well management routine computes the current state of surface facilities and wells from the governing equations. At block 406, the values from the well management routine are used along with the value of state variables at each computational cell to construct a Jacobian matrix. The Jacobian matrix is the matrix (or array) of all first order partial derivatives (with respect to the state variables) of a vector valued function. In reservoir simulation, the Jacobian details the change of the governing partial differential equations with respect to the state variables (pressure, saturation).

[0058] At block 408, the linear solver uses the Jacobian matrix to generate updates for physical properties of interest, such as pressure and saturation, among others. At block 410, the calculated physical properties are compared to either previously calculated properties or to measured properties, and, at block 412, a determination is made as to whether a desired accuracy has been reached. In an exemplary embodiment, the determination is made by determining that the calculated properties have not significantly changed since the last iteration (which may indicate convergence). For example, convergence may be indicated if the currently calculated properties are within 0.01%, 0.1%, 1 %, 10 %, or more of the previously calculated properties. In other embodiments, the determination may be determining if the calculated properties are sufficiently close to measured properties, for example, within 0.01%, 0.1%, 1 %, 10 %, or more. If the desired accuracy is not reached, process flow returns to block 408 to perform another iteration of the linear solver.

[0059] If at block 412, the desired accuracy has been reached, process flow proceeds to block 414, at which results are generated. The results may be stored in a data structure on a

Attorney No.: 2010EM054

tangible, machine readable medium, such as a database, for later presentation, or the results may be immediately displayed or printed after generation. At block 416, the time is incremented by a desired time step, for example, a day, a week, a month, a year, 5 years, 10 years or more, depending, at least in part, on the desired length of time for the simulation. At
5 block 418, the new time is compared to the length desired for the simulation. If the simulation has reached the desired length of time, the simulation ends at block 420. If the time has not reached the desired length, flow returns to block 404 to continue with the next increment.

[0060] The parallelization of the processes may be considered to fall into two main types,
10 task based parallelization and grid based parallelization. For task based parallelization, a calculation is divided into sub tasks that are run independently in parallel. For example, in the well management task at block 404, a set of operations may be computed on each of a set of wells that can be performed independently of one another. Therefore each computing unit may execute the operations independently of the other computing units unit.

[0061] Grid based parallelization may be performed at a number of points in the
15 processes, such as in the Jacobian construction and/or the property calculations discussed with respect to blocks 406 and 410. In the computational process of constructing the Jacobian, rock and fluid properties with corresponding derivatives are calculated at each computational cell. This type of parallelization is used for computations that do not depend
20 on the computational cells being adjacent or require global communication for the computations.

[0062] Vapor-liquid equilibrium (VLE) fluid property computations may be considered
in an example of parallelization. If a simulation model uses a black oil fluid characterization for VLE computations, the amount of computational work required for a flash calculation is
25 roughly proportional to the number of computational cells due to the linear nature of the black oil VLE computations. However, if a compositional fluid model is chosen, the amount of computational work for the flash calculation within a single computational cell depends on the cell's position in the phase state space. Hence, the amount of computational work may vary sharply from cell to cell.

[0063] Grid based parallelization may also be used to divide the problems used by the
30 linear solver among computing units. The linear solver is an algorithm used to compute

Attorney No.: 2010EM054

updates to state properties (pressure, saturation, etc). A linear solver requires local computations that are similar between subdomains, in other words the computational cells associated with each computing unit. However, linear solvers also require global computations (which are not required for Jacobian construction) to compute updates to the
5 state variables of interest.

[0064] Partitioning a simulation model generally has three goals: load balancing, communication minimization, and algebraic compatibility. A good load balance evenly distributes the computational work from one computing unit to another. Further, a good partitioning minimizes the interfacing used between partitions, thus minimizing the
10 communication costs required by the algorithm using the partition. A good partitioning also asserts algebraic compatibility, allowing fast convergence of the linear solver. Efficient partitioning may decrease processing time at all stages of a simulation model, but is especially useful for the efficiency of a parallel linear solver.

[0065] As previously noted, one of the factors that can influence the efficiency of a
15 parallel linear solver is assigning computational cells representing a well bore (and near well bore region) to a single computing unit. Therefore, the boundaries of the subdomains should not partition the well connections otherwise the convergence of a linear solver may dramatically degrade.

[0066] Partitioning can be based on a topological graph of the matrix connectivity. In an
20 exemplary embodiment, the topological graph is generated from a connectivity diagram of the computational mesh wherein the center of each computational cell represents a node, and a graph edge connect each node (or center point). For the computational mesh, as discussed with respect to Fig. 3, the nodes of the graph corresponding to adjacent computational cells are connected by graph edges.

[0067] In another exemplary embodiment, the topological map is generated from a
25 connectivity diagram of the matrix representation of the linear system. For the matrix, the diagonal elements of the matrix represent the nodes of the graph and the non-zero off-diagonal matrix elements denote a connection between corresponding nodes. In addition to the connectivity diagram, partitioning algorithms may consider the weights assigned to edges
30 of the graph (connections). Typically, a larger connection weight implies a smaller chance

Attorney No.: 2010EM054

that particular connection will be cut and the nodes to which this edge is connected will be partitioned into separate subdomains by the partitioning algorithm.

Partitioning

[0068] An exemplary embodiment of the present techniques defines a weighting scheme for the connections in a topology map that are perforated by a wellbore so that a partitioning algorithm (for example, METIS or Chaco) preserves wells, near well nodes, and reservoir cells connected to those wells in a single subdomain. The weighting scheme is not limited to wells, but can also be applied to any reservoir feature that should remain in a single subdomain when partitioning.

10 [0069] Fig. 5 is a block diagram of a method 500 for partitioning a simulation model, in accordance with an exemplary embodiment of the current techniques. The method 500 begins at block 502 with the construction of a topology graph of grid connectivity. The topological map may be created from the computational mesh discussed with respect to Fig. 2, for example, by associating a center point with each computational cell and then
15 connecting each center point, as discussed with respect to Fig. 3. The connections between the center points form the topological graph of grid connectivity. At block 504, initial weights for the connections may be created using physical information (such as transmissibility or total velocity, among others) for each connection on the topological graph.

[0070] At block 506, a mapping or table of connections between wells and grid
20 connections (or reservoir nodes) can be created. The mapping may be used to create large connection weights for connections that are perforated by the same well or reservoir feature.

[0071] If the values used for scaling the weights are very large, non-integer, or both, it may be necessary to scale and map the physical information to integer graph weights, as determined at block 508. The scaling is performed at block 510, and may, for example, use
25 the techniques discussed below with respect to Figs. 7 and 8. In some embodiments, a threshold cut-off may be applied to the weights. For example, if a connection weight exceeds an upper limit, the weight may be set to the upper limit.

[0072] At block 512, a partitioning algorithm may use the weighted graph as an input. The partitioning algorithm divides the computational mesh into subdomains, based on the
30 weighting of the connections. Higher weighted connections are less likely to be divided by partitions into separate subdomains than lower weighted connections. Accordingly, the

Attorney No.: 2010EM054

partitions between subdomains will follow local minima in the topology graph, ensuring that wells and other significant reservoir features are kept in single subdomains. At block 514, the output of the partitioning algorithm is used for the assignment of computational cells to processors for the linear solver (or for assigning other portions of the simulation model).

5 [0073] Fig. 6 is a block diagram of a method 600 for partitioning a simulation model, in accordance with an exemplary embodiment of the present techniques. At block 602, a topology graph is created from the linear system (or Jacobian) matrix that is to be partitioned. A graph is a collection of nodes (or vertices) and a collection of edges (or connections) that connect pairs of nodes. Weights can be assigned to the edges to form a weighted graph. A
10 graph can be created from a linear system matrix by representing each linear system row (or equation) as a node and the connection between nodes as an edge. For example, the connections in the topology graph can be created for off-diagonal entries in a Jacobian matrix representing the well simulation.

[0074] At block 604, initial weights for the graph connections may be created using the
15 values of the off-diagonal matrix entries corresponding to each connection on the topological graph. At block 606, a mapping or table of connections between wells and reservoir nodes (i.e., which matrix equation unknowns correspond to the wells and to which reservoir unknowns they are connected) can be created. As for the physical mapping discussed with respect to Fig. 5, this mapping can be used to create larger connection weights between the
20 nodes of the graph corresponding to the connections that are perforated by the same well.

[0075] At block 608, a determination is made as to whether scaling is needed, for example, if the connection weights are too large for use by the partitioning algorithm. At block 610, the matrix values are scaled to the integer graph weights, for example, by doubling the matrix values. If necessary, a threshold cut-off is applied to the weights. The scaling is
25 discussed further with respect to Figs. 7 and 8.

[0076] At block 612, a partitioning algorithm may use the weighted graph as an input. The partitioning algorithm divides the computational mesh into subdomains, based on the weighting of the connections. As discussed above, higher weighted connections are less likely to be divided by partitions into separate subdomains than lower weighted connections.
30 At block 614, the subdomains are used as a domain mapping for the assignment of grid blocks to processors for the linear solver.

Attorney No.: 2010EM054

[0077] In an exemplary embodiment, the method discussed with respect to Figs. 5 and 6 is applied once during a simulation to create a static partition. In other embodiments, the partitioning method can be used to create a dynamic partition by applying the method at different stages of the simulation using changing data (for example, total velocity) to define
5 the weights of the graph or matrix.

Scaling Connection Weights Between Cells

[0078] Exemplary embodiments of the present techniques use special values for the weights associated with wells and map those weights into integer values for input into the partitioning algorithms. However, the range of physical values may be extremely large, and,
10 thus, the weights may need to be scaled. For example, if transmissibility coefficients are used for initial weight definition, the range can easily span from 10^{-6} to 10^2 Darcy. Further, partitioning algorithms, such as METIS or Chaco, typically accept weights as positive integer values within a relatively small range, for example, between 1 and 1000. A larger range of integer values can produce a less robust partition. Accordingly, a proper mapping between
15 the large range of physical weights and (relatively) small range of integer weights can be useful. For example, a scaling or mapping scheme can be a monotone, linear mapping that filters out extreme values of the connection weights.

[0079] As described below, the weighting of connections may be performed by several techniques. Specifically, the weighting may be performed on the basis of perforations, near
20 well (or feature) regions, and flow calculations, among others. Each of these techniques is described below.

[0080] In an exemplary embodiment of the present techniques, a table or list of grid blocks or matrix rows that are perforated by the well bore is created. The grid blocks or matrix rows that appear in this list or table are used for the weighting scheme. For example,
25 the connection weights may have an incremental value added to the physical weighting for each connection that is perforated by a well.

[0081] In another exemplary embodiment, a near well procedure may be used. The near well procedure determines which grid blocks or matrix rows are approximate to the well perforations, for example, within two or three adjacent blocks or rows, and includes those
30 grid blocks or matrix rows in the weighting scheme. For example, a first value may be added to the weight for each connection that is perforated by a well, and a second value may be

Attorney No.: 2010EM054

added to the weight for each connection that is adjacent to a connection that is perforated by a well.

[0082] Another technique that may be used in exemplary embodiments chooses which grid blocks to include in the weighting scheme by solving a local flow problem to determine how the fluids travel away from the well bore. The results of this flow problem tell which
5 grid blocks or matrix rows should remain in the same subdomain, allowing those grid blocks or matrix rows to have increased weights, for example, by adding a value.

[0083] The techniques for scaling the maps are not limited to those discussed above. One of ordinary skill will recognize that any number of other techniques may be used to weight
10 the connections. For example, statistical calculations or artificial intelligence systems (such as neural networks), may be used to determine the weighting for connections near wells and features. All of these techniques follow the method shown in Fig. 7.

[0084] Fig. 7 is a process flow diagram illustrating a method 700 for scaling connection weights, in accordance with an exemplary embodiment of the present techniques. The method generally corresponds to block 510 or block 610, discussed with respect to Figs. 5
15 and 6, respectively. The method 700 begins at block 702 with the scaling of the range of real weight values to the interval $[0,1]$. This may be performed by a direct linear mapping or by a statistical calculation. For example, if connections in the topology graph are weighted using transmissibility, a cumulative distribution function (or similar probability distribution
20 function) which describes the distribution of the transmissibility value that can be used to scale the transmissibility values into the interval $[0,1]$.

[0085] At block 704, scaled weights that are within a defined threshold range of 1 are reset to the bottom of the scaled range. For example, for some small threshold value ε , connection weights within the interval $[1-\varepsilon,1]$ can be reassigned to a weight of $1-\varepsilon$.
25 Typically, a threshold value of $\varepsilon = 0.10$ or $\varepsilon = 0.05$ is selected and larger weights are assigned (such as transmissibility values) to the value: $1-\varepsilon$. In an exemplary embodiment, the value for $\varepsilon = 0.05$

[0086] Once we truncate the large values, at block 706, a linear or nonlinear monotone mapping function can be used to map the interval $[0, 1-\varepsilon]$ to the integer values between 1 and
30 an upper integer value that is consistent with the partitioning algorithm input. The upper

Attorney No.: 2010EM054

integer value may be 100, 500, 1000, 10000, or more. In other words, the graph connection weights \tilde{T} can be mapped to integer values W in a predefined range, typically from 1 to $N=1000$ or 10,000.

[0087] Fig. 8 is a graph 800 illustrating a mapping of real values to integer weights using a cumulative distribution function, in accordance with an exemplary embodiment of the present techniques. In the graph 800, the x-axis 802 represents transmissibility in the reservoir, while the y-axis 804 represents the value of the cumulative distribution function (CDF). The interval $[0, 1-\varepsilon]$ is indicated by reference number 806. The CDF_{cutoff} 808 is the value $1-\varepsilon$. Thus, any value for T that generates a mapped value for the CDF that is greater than the CDF_{cutoff} 808 is reassigned to the value of T at the CDF_{cutoff} 808. The value of T at the CDF_{cutoff} 808 is termed the T_{cutoff} 810.

[0088] Once the values for T are assigned, the integer connection weights W can be calculated by the formulas shown in Eqn. 1.

$$\tilde{T} = \begin{cases} T, & T < T_{cutoff} \\ T_{cutoff}, & T \geq T_{cutoff} \end{cases}, \quad W = 1 + (N - 1) \cdot [\tilde{T} / T_{cutoff}] \quad \text{Eqn. 1}$$

As indicated in Eqn. 1, the non-integer connection weights \tilde{T} for transmissibility may be set to the value of transmissibility, if less than T_{cutoff} 810, and set to T_{cutoff} 810 if equal to or greater than T_{cutoff} 810. After the non-integer connection weights are computed, the integer connection weights, W , may be calculated using the formula shown in Eqn. 1. The procedure for mapping the connection weights to an integer range is not limited to the formula shown in Eqn. 1. In other embodiments, W may be calculated any other monotone (or order preserving) mapping algorithm, such as the formula shown in Eqn. 2.

$$W = 1 + (N - 1) \cdot [\ln(\tilde{T} / T_{cutoff}) / \ln(2)] \quad \text{Eqn. 2}$$

[0089] The combination of scaling with a cut-off threshold and a proper monotone mapping constructs connection weights that allow a partitioning algorithm to partition a reservoir without placing a well or other reservoir feature into different subdomains. The weighting also keeps the nodes and edges corresponding to the large values of the physical property used to assign the original real weights (such as transmissibility or total velocity) in a single subdomain, which may enhance the performance of parallel linear solver.

Attorney No.: 2010EM054

Exemplary Cluster Computing System

[0090] Fig. 9 is a block diagram of an exemplary cluster computing system 900 that may be used in exemplary embodiments of the present techniques. The cluster computing system 900 illustrated has four computing units 902, each of which may perform calculations for part of the simulation model. However, one of ordinary skill in the art will recognize that the present techniques are not limited to this configuration, as any number of computing configurations may be selected. For example, a small simulation model may be run on a single computing unit 902, such as a workstation, while a large simulation model may be run on a cluster computing system 900 having 10, 100, 1000, or even more computing units 902. In an exemplary embodiment, each of the computing units 902 will run the simulation for a single subdomain. However, allocation of the computing units 902 may be performed in any number of ways. For example, multiple subdomains may be allocated to a single computing unit 902 or multiple computing units 902 may be assigned to a single subdomain, depending on the computational load on each computing unit 902.

[0091] The cluster computing system 900 may be accessed from one or more client systems 904 over a network 906, for example, through a high speed network interface 908. Each of the client systems 904 may have tangible, computer readable memory 910 for the storage of operating code and programs, including random access memory (RAM) and read only memory (ROM). The operating code and programs may include the code used to implement all or portions of the methods discussed with respect to Figs. 4-7. The client systems 904 can also have other tangible, computer readable media, such as storage systems 912. The storage systems 912 may include one or more hard drives, one or more optical drives, one or more flash drives, any combinations of these units, or any other suitable storage device. The storage systems 912 may be used for the storage of code, models, data, and other information used for implementing the methods described herein.

[0092] The high speed network interface 908 may be coupled to one or more communications busses in the cluster computing system 900, such as a communications bus 914. The communication bus 914 may be used to communicate instructions and data from the high speed network interface 908 to a cluster storage 916 and to each of the computing units 902 in the cluster computing system 900. The communications bus 914 may also be used for communications among computing units 902 and the storage array 916. In addition

Attorney No.: 2010EM054

to the communications bus 914 a high speed bus 918 can be present to increase the communications rate between the computing units 902 and/or the cluster storage 916.

5 [0093] The cluster storage 916 can have one or more tangible, computer readable media devices, such as storage arrays 920 for the storage of data, visual representations, results, code, or other information, for example, concerning the implementation of and results from the methods of Figs. 4-7. The storage arrays 920 may include any combinations of hard drives, optical drives, flash drives, holographic storage arrays, or any other suitable devices.

10 [0094] Each of the computing units 902 can have a processor 922 and associated local tangible, computer readable media, such as memory 924 and storage 926. The memory 924 may include ROM and/or RAM used to store code, for example, used to direct the processor 922 to implement the methods illustrated in Figs. 4-7. The storage 926 may include one or more hard drives, one or more optical drives, one or more flash drives, or any combinations thereof. The storage 926 may be used to provide storage for intermediate results, data, images, or code associated with operations, including code used to implement the methods of
15 Figs. 4-7.

[0095] The present techniques are not limited to the architecture of the cluster computer system 900 illustrated in Fig. 9. For example, any suitable processor-based device may be utilized for implementing all or a portion of embodiments of the present techniques, including without limitation personal computers, laptop computers, computer workstations, GPUs,
20 mobile devices, and multi-processor servers or workstations with (or without) shared memory. Moreover, embodiments may be implemented on application specific integrated circuits (ASICs) or very large scale integrated (VLSI) circuits. In fact, persons of ordinary skill in the art may utilize any number of suitable structures capable of executing logical operations according to the embodiments.

25 [0096] While the present techniques may be susceptible to various modifications and alternative forms, the exemplary embodiments discussed above have been shown only by way of example. However, it should again be understood that the present techniques are not intended to be limited to the particular embodiments disclosed herein. Indeed, the present techniques include all alternatives, modifications, and equivalents falling within the true spirit
30 and scope of the appended claims.

Attorney No.: 2010EM054

CLAIMS

1. A method for partitioning a reservoir, comprising:
generating a representation of a topology graph of a simulation model in a tangible,
computer readable medium, wherein the topology graph comprises a plurality
5 of elements and a plurality of connections between adjacent elements;
weighting each of the plurality of connections to create a plurality of weights;
scaling each of the plurality of weights;
partitioning the topology graph into two or more subdomains, wherein a partition
boundary follows a local topographical minimum in the topology graph; and
10 assigning a subdomain to each of a plurality of processors.
2. The method of claim 1, wherein the plurality of elements comprise
computational cells in a computational mesh.
- 15 3. The method of claim 1, wherein the plurality of elements comprise rows in a
linear system matrix and the plurality of connections correspond to nonzero elements of the
matrix.
- 20 4. The method of claim 2, wherein the weighting of each of the plurality of
connections is based, at least in part, on physical properties assigned to the computational
cells.
- 25 5. The method of claim 4, wherein the physical properties comprise
transmissibility, total mobility, mass flow, heat flow, or any combinations thereof.
6. The method of claim 3, wherein the weighting of each of the plurality of
connections is based at least in part on a plurality of off-diagonal coefficients in a Jacobian
matrix representing a simulation model.

Attorney No.: 2010EM054

7. The method of claim 1, wherein the weighting of each of the plurality of connections is based, at least in part, upon its proximity to a well bore.

5 8. The method of claim 1, wherein a near-well region is kept in one of the two or more subdomains.

9. The method of claim 1, wherein the weighting of each of the plurality of connections is based, at least in part, upon its belonging to a grid block that is perforated by a
10 well bore.

10. The method of claim 1, wherein the weighting of each of the plurality of connections is based, at least in part, upon a solution to a local flow problem.

15 11. The method of claim 1, wherein scaling each of the plurality of weights is performed using a probability distribution.

12. The method of claim 1, wherein scaling each of the plurality of weights is performed by scaling discrete reservoir properties to create a linear mapping.

20

13. The method of claim 1, further comprising performing the simulation model for a subdomain assigned to each of the plurality of processors.

14. A system for modeling reservoir properties, comprising:

25 a plurality of processors;

a storage medium comprising a representation of a topology graph of a reservoir model, wherein the topology graph comprises a plurality of connections; and

Attorney No.: 2010EM054

a machine readable medium comprising code configured to direct at least one of the plurality of processors to:

weight each of the plurality of connections to create a plurality of weights;

map each of the plurality of weights to an integer value;

5 partition the topology graph into two or more subdomains; and

assign each of the two or more subdomains to one of the plurality of processors.

15 15. The system of claim 14, wherein each of the plurality of weights is based, at least in part, on an associated physical property.

16. The system of claim 15, wherein the physical property represents heat transfer, mass transfer, total flow, transmissibility, or any combinations thereof.

15 17. The system of claim 14, wherein the plurality of processors comprises a cluster computing system.

18. A tangible, computer readable medium comprising code configured to direct a processor to:

20 create a topology graph of a simulation model, wherein the topology graph comprises a plurality of connections between the center points of adjacent computational cells in a computational mesh;

weight each of the plurality of connections to create a plurality of weights;

scale each of the plurality of weights;

25 partition the topology graph into two or more subdomains, wherein a partition boundary follows a local topographical minimum in the topology graph; and


assign each of the subdomains to one of a plurality of processors.

Attorney No.: 2010EM054

19. The tangible, computer readable medium of claim 18, comprising code configured to direct the one of the plurality of processors to process a simulation model for the subdomain.

5

20. The tangible, computer readable medium of claim 18, comprising code configured to direct the processor to map the plurality of weights into an integer range.



100
FIG. 1





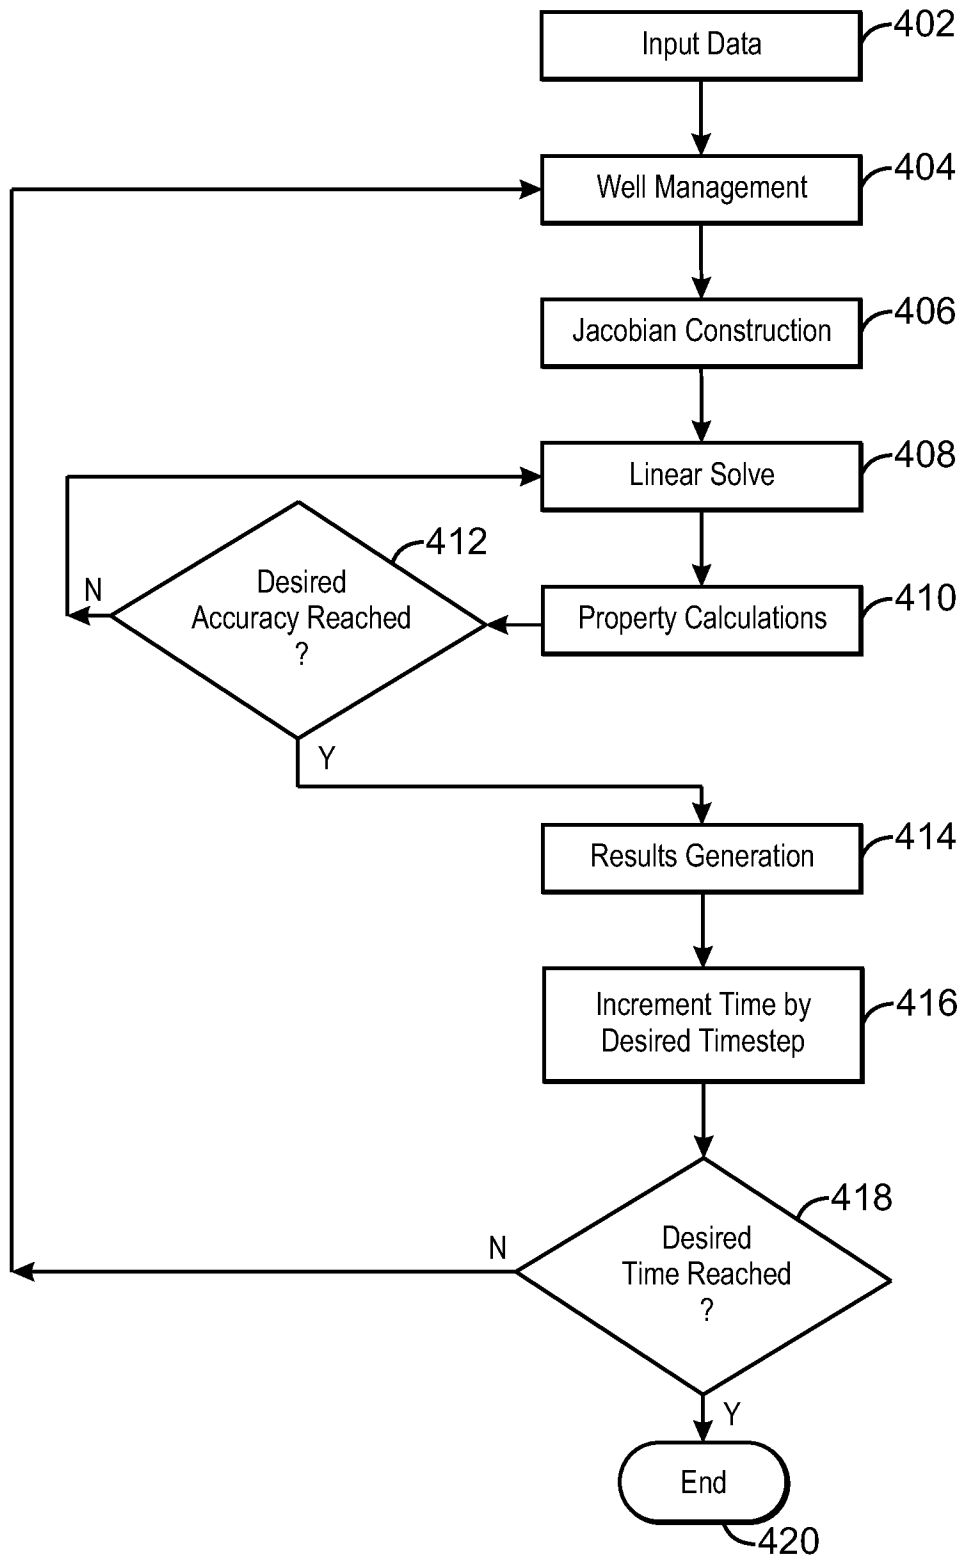
FIG. 2

3/9



300
FIG. 3

4/9



400
FIG. 4

5/9

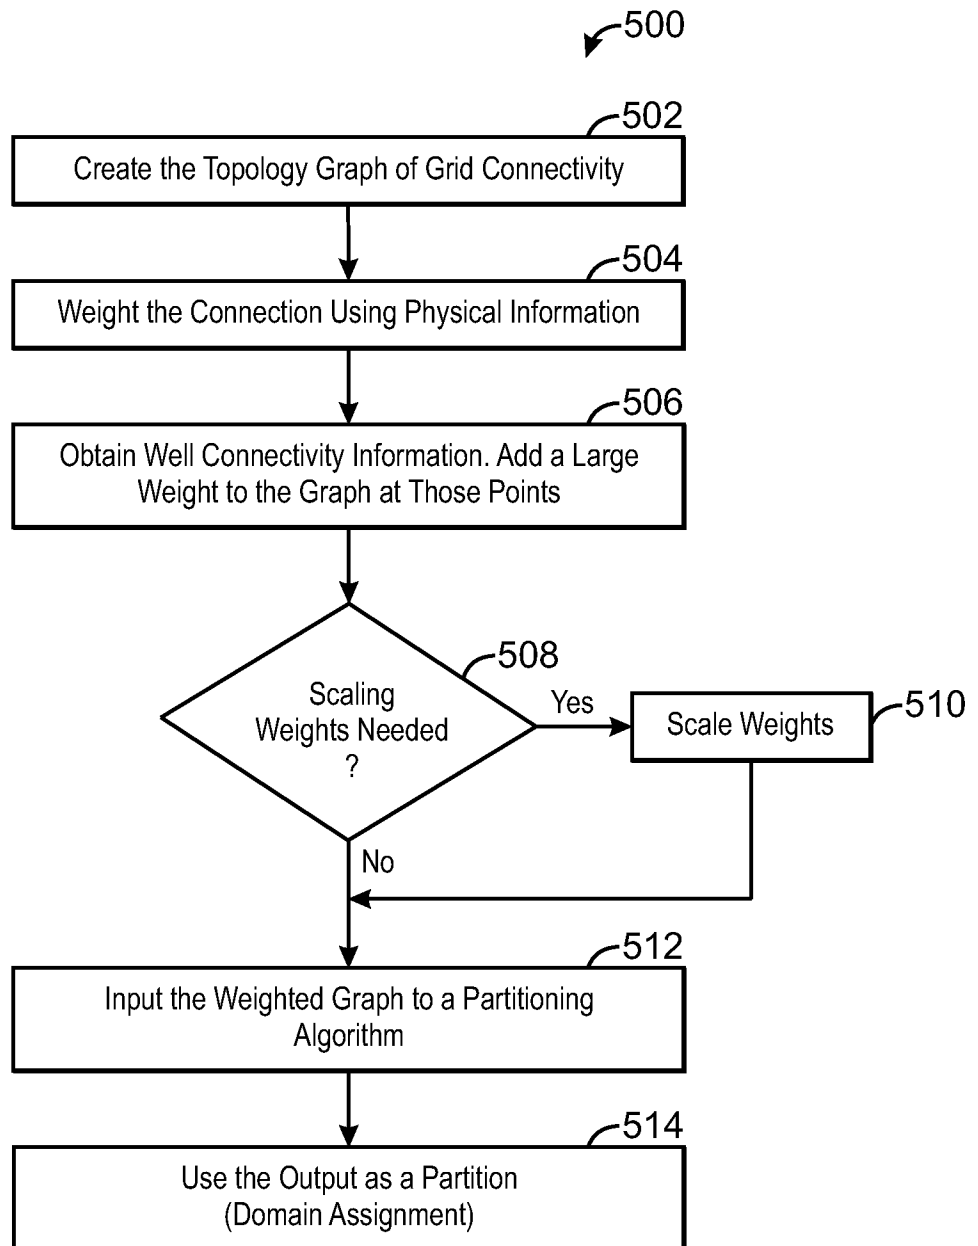


FIG. 5

6/9

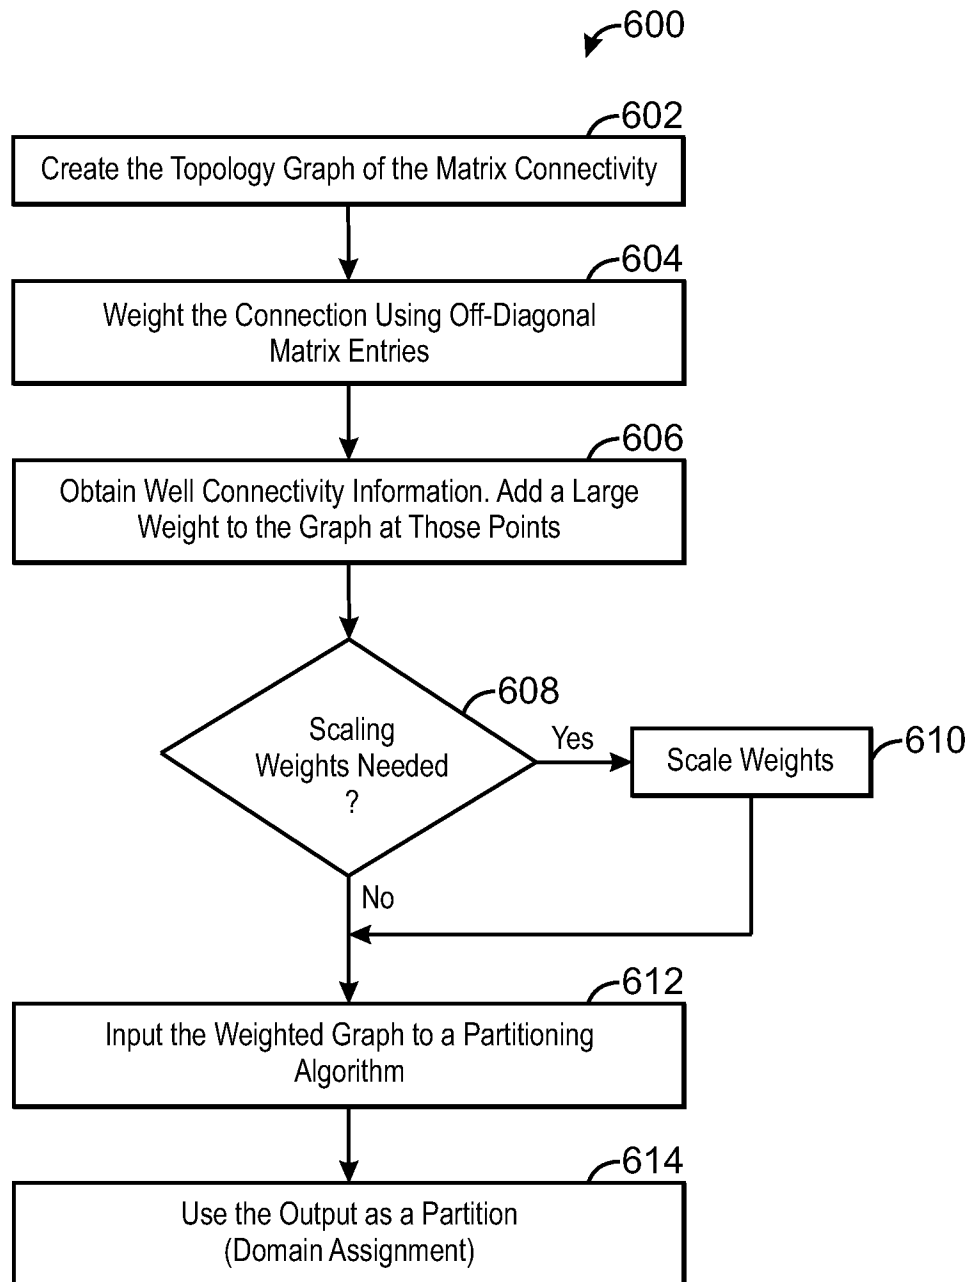


FIG. 6

7/9

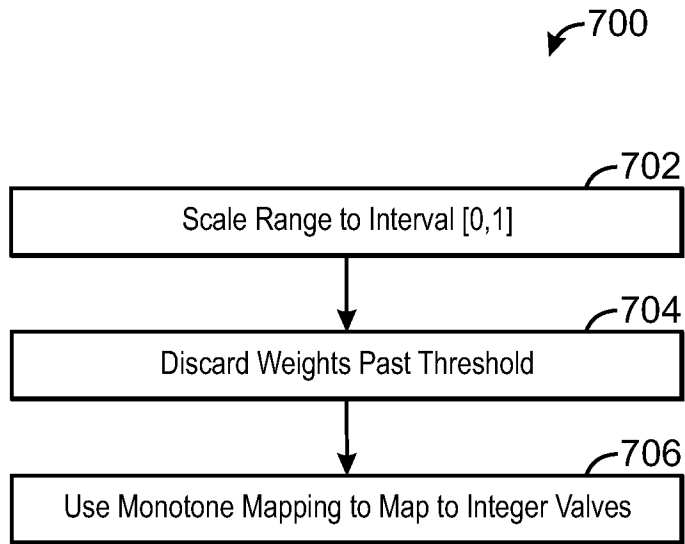


FIG. 7

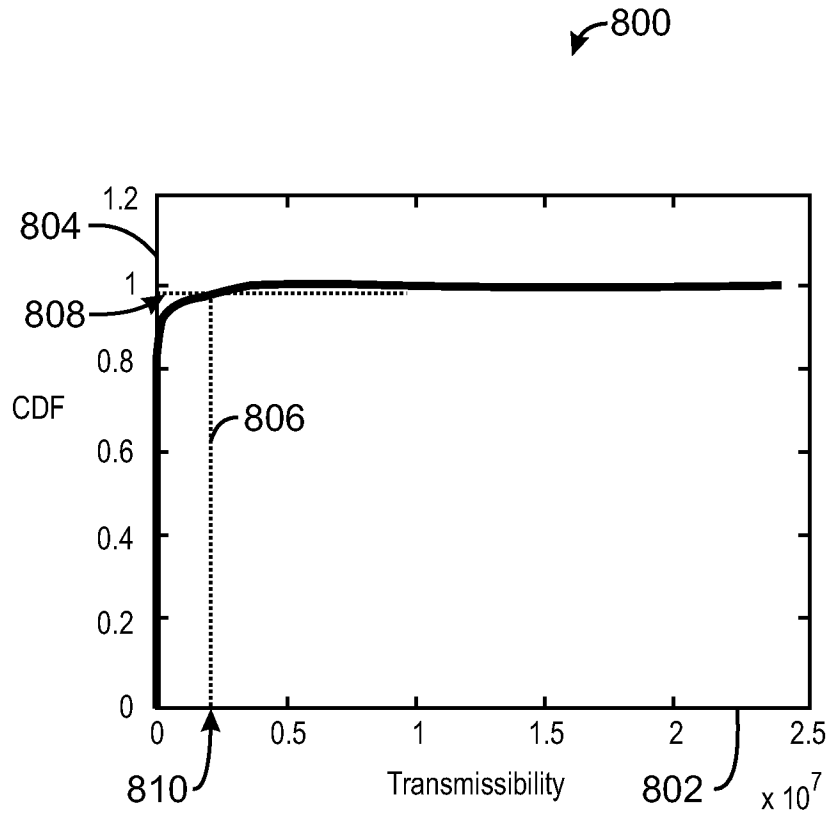


FIG. 8




FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 10/53141

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06G 7/48 (2010.01)

USPC - 703/10

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8): G06G 7/48 (2010.01)

USPC: 703/10

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 703/6,17,22 (keyword limited; terms below)

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

PubWest; Google Scholar; Google Patents; FreePatentsOnline. Search terms used:

simulation emulation forecasting prediction model, reservoir well well-bore, parallel simultaneous same-time, processor microprocessor
CPU SPU ASIC DSP computation cluster, partition segment section subset domain sub-domain cell ...

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2009/0248374 A1 (HUANG et al.) 01 October 2009 (01.10.2009) entire document, especially Abstract; Fig.3; para [0078], [0079], [0089], [0129], [0140], [0146], [0157], [0183], [0191], [0193], [0214], [0224]	1 - 20
Y	US 2002/0013687 A1 (ORTOLEVA) 31 January 2002 (31.01.2002) entire document, especially Abstract; para [0015], [0206], [0239], [0286], [0292], [0296]	1 - 20
Y	US 2006/0025976 A1 (KENNON et al.) 02 February 2006 (02.02.2006) entire document, especially Abstract; Fig.13; para [0042], [0082], [0095], [0120], [0128]	6, 7, 12
Y	US 2006/0235667 A1 (FUNG et al.) 19 October 2006 (19.10.2006) entire document, especially Abstract; para [0019], [0055]	17
A	US 2005/0015231 A1 (EDWARDS et al.) 20 January 2005 (20.01.2005) entire document	1 - 20
A	US 2007/0027666 A1 (FRANKEL) 01 February 2007 (01.02.2007) entire document	1 - 20

 Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

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

"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

06 December 2010 (06.12.2010)

Date of mailing of the international search report

16 DEC 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774