

(43) International Publication Date
23 August 2012 (23.08.2012)(51) International Patent Classification:
G06F 11/14 (2006.01) *G06F 9/48 (2006.01)*

(74) Agents: HENNESSEY, Gilbert H. et al.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (US).

(21) International Application Number:

PCT/US2012/025388

(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, QA, RO, RS, RU, RW, 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.

(22) International Filing Date:
16 February 2012 (16.02.2012)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/030,998 18 February 2011 (18.02.2011) US

(71) Applicant (for all designated States except US): AB INITIO TECHNOLOGY LLC [US/US]; 201 Spring Street, Lexington, Massachusetts 02421 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): DOUROS, Bryan Phil [US/US]; 92 Lakeview Road, Framingham, Massachusetts 10701 (US). WHOLEY, Joseph Skeffington III [US/US]; 11 Hillcrest Road, Belmont, Massachusetts 02478 (US).

(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, 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).

[Continued on next page]

(54) Title: RESTARTING PROCESSES

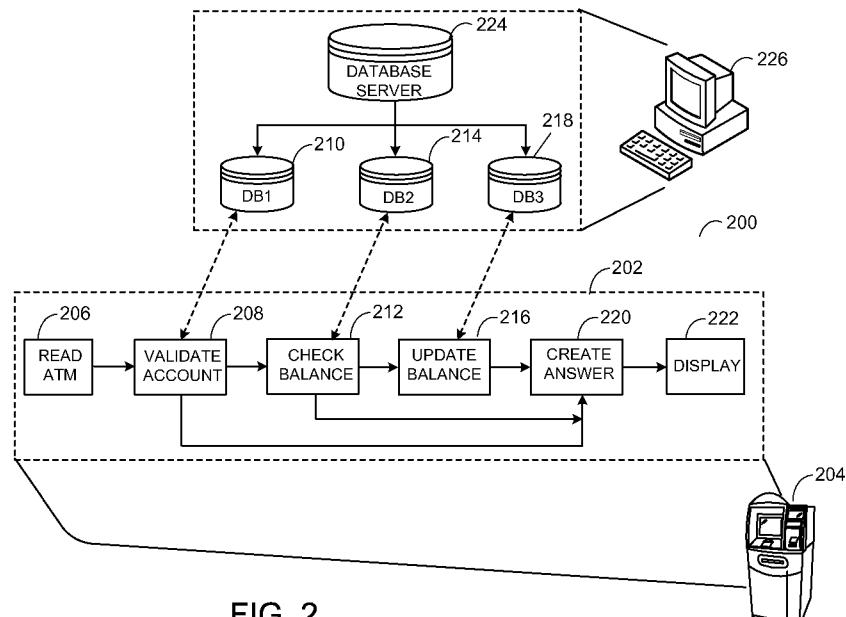


FIG. 2

(57) **Abstract:** Techniques are disclosed that include a computer-implemented method, including storing information related to an initial state (402) of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state (404) of the execution phase; aborting execution (506) of the process in response to a predetermined event; and resuming execution of the process from one of the saved initial and end states (512) without needing to shut down the process.



Published:

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

RESTARTING PROCESSES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Application Serial No. 13/030,998, filed on February 18, 2011, entitled “Restarting Processes”, the entire contents of which are hereby incorporated by reference.

BACKGROUND

This description relates to restarting processes.

Computational speeds provided by single processor computers have advanced tremendously over the past decades. However, many applications executed by such processors may need computational capacities that exceed even the fastest single processor computer. For example, in a transactional system, e.g., an airline reservation system, multiple users may concurrently access computer resources. These users typically expect low response times. A single process computer may not be able to keep up with such demand. A variety of architectures such as parallel processing systems have been developed to handle such applications to improve performance. In general, parallel processing systems use multiple processors that may be located at a single site or remotely distributed. Due to their processing capabilities, such parallel processing systems have become relied upon for applications that process large volumes of data, which in some cases can include essentially continuous and near real-time processing. Such processing capabilities are expected to be robust and resistant to system failures, i.e., fault tolerant. These capabilities are useful for all kinds and sizes of computer networks ranging from large-scale Internet-based data processing to private networks and communication systems (e.g., internal corporate “intranets”, etc.).

SUMMARY

In one aspect, in general, a computer-implemented method includes storing information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of

an end state of the execution phase; aborting execution of the process in response to a predetermined event; and resuming execution of the process from one of the saved initial and end states without needing to shut down the process.

Aspects can include one or more of the following.

The predetermined event can represent loss of connection to an external device. The predetermined event represents an error with an external device. The execution of the process can be resumed when the connection to the external device has been restored. The execution of the process can be resumed when the error with the external device has been cleared. The execution of the process can be resumed from an end state that is stored prior to an execution phase in which the predetermined event occurred. The execution of the process can be resumed from the initial state if the predetermined event occurs during an execution phase substantially immediately after startup of the process. The execution of an execution phase can include performing one or more processing actions on a received stream of data to produce output data corresponding to the execution phase.

The method can further include storing the output data corresponding to one or more execution phases, and reproducing the output data when execution of the process is resumed. The process may be part of a process stage and in communication through a data path to a second, different process in the process stage. The method can further include passing a checkpoint command message through each process in the process stage; and at each process, saving new information related to an initial or end state upon receiving the checkpoint command message, where the saving includes suspending operation of the process and saving the new information to a storage area. The computer-implemented method can further include overwriting an old saved initial or end state with the new initial or end state.

Each process may be in communication with a data queue for receiving and queuing data for the process. The method can further include generating the checkpoint command message in response to detecting a trigger event. The trigger event can include information about a network event. The method can further include periodically generating the checkpoint command message.

The method can further include generating the checkpoint command message in response to an occurrence of selected data values within or derived from incoming data records being processed. The method can further include generating an abort command message in response to the predetermined event; passing the abort command message through each process in the process stage; and on receiving the abort command message at each process, aborting execution of the process and passing the abort command message to a next process. The method can further include resuming execution of the process from a selected one of the saved initial or end states based in part on information contained in a resume processing message. The method can further include receiving the abort command message during a first execution phase substantially immediately after the process is initialized and information related to the initial state of the process has been saved; and resuming execution of the process from the saved initial state without need to shut down and restart the process.

In another aspect, in general, a computer-readable storage medium storing a computer program includes instructions for causing a computing system to store information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase; abort execution of the process in response to a predetermined event; and resume execution of the process from one of the saved initial and end states without needing to shut down the process.

In another aspect, in general, a computing system includes an input device or port configured to receive and store information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase; and at least one processor configured to: abort execution of the process in response to a predetermined event, and resume execution of the process from one of the saved initial and end states without needing to shut down the process.

In another aspect, a computing system includes means for storing information related to an initial state of a process upon being initialized, wherein execution of the

process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase; means for controlling execution of the process, the controlling including causing the aborting of execution of the process in response to a predetermined event, and resuming execution of the process from one of the saved initial and end states without needing to shut down the process.

Aspects can include one or more of the following advantages.

Processes in multi-process processing systems can be executed in distinct execution phases. In the event of a system failure, terminating and restarting a processing system from the most recently complete checkpoint can consume an undue amount of processing time and resources. After a processing system has terminated its activity in response to the exception condition, the processing system may need to be manually reinitialized by an information technology specialist experienced with such systems. This can result in significant system downtime. In some examples, a separate process may also be needed to detect the system failure and inform a specialist. As such, to improve efficiency and reduce processing resource consumption, processes within the processing system may be executed from their last recorded checkpoints instead of restarting the entire system upon a failed connection to a process being restored. In one implementation, rather than terminating and restarting the entire system, the individual processes in the system may be informed to suspend processing until the failed connection is restored.

Other features and advantages of the invention will become apparent from the following description, and from the claims.

DESCRIPTION

FIG. 1 is a block diagram of a multi-process data processing system.

FIGS. 2 and 3 illustrate exemplary multi-process data processing systems.

FIG. 4 is flowchart illustrating an exemplary checkpointing process.

FIGS. 5 and 6 are flowcharts of exemplary recovery mechanisms.

Referring to FIG. 1, a data processing system 100 provides multiple processes arranged in a streamlined manner for processing data. Within the exemplary system 100,

data is received from a data source 102 (e.g., an application being executed on a server 104 that functions as a Web server) and communicated to a multi-process data processing module 106 being executed on a computer system 108 or executed in a distributed manner (e.g., with two or more networked computer terminals). The data processing module 106 monitors, controls, and performs the data processing aspects of the system 100. To provide such processing, the data processing module 106 includes one or more queues 110, 112, 114 capable of storing data to be processed by one or more processes 116, 118. In this instance, as shown, data received from the data source 102 is stored in an initial data queue 110 and periodically provided to the initial process 116. The process 116 processes the data (e.g., transforms, filters, confirms content, etc.) and provides the processed data to one or more downstream data queues 112, 112'. Subsequent processes 118, 118' may be provided data from the queues 112, 112' and perform other (or similar) processing before in turn delivering results to other downstream data queues 114, 114'. The illustrated queue and process layout of data processing module 106 is one of many possible processing schemes that may be utilized. For example, the data processing module 106 may include additional processes (e.g., for parallel or serial execution) that may be located upstream, downstream or independent of the shown processes. In some examples, data from the last set of queues (e.g., queues 114 and 114') may be output to a destination application 120 (or multiple applications), such as a relational database management system (RDBMS).

The processes included in the data processing module 106 may be in communication with external devices and/or other processing systems (e.g., a computer system 122). For example, the processes may be in communication with a Java Message Service (JMS) queue that provides messages to the processes from other systems. In some cases, the processes in the data processing module 106 may be in communication with one or more databases (e.g., located in the external system 122). For example, the module 106 may perform updates to customers' financial accounts in a bank database based on information received from corresponding customer sessions at one or more automated teller machines (ATM).

By way of example, FIG. 2 illustrates a processing system 200 having a remotely executed processing module 202 (executed by an ATM 204) being used to provide data

for processing at a central location. In the illustrated example, an initial process, e.g., a read ATM process 206 is capable of receiving customer account data (e.g., associated with a transaction) from an ATM and passing the data to a validate account process 208 for authenticating the account details. In this instance, the validate account process 208 can verify a personal identification number (PIN) entered by the customer against a PIN database 210. Once the customer's identity has been authenticated, further data records may be communicated downstream to a check balance process 212, which may communicate with a second, different database 214, e.g., for checking the balance of the identified customer account. After further transactions are completed, additional data may be sent downstream to an update balance process 216 which may communicate with a third database 218, e.g., for updating the balance information associated with the customer account. A create answer process 220 may prepare an output summary of the transactions, which may be provided to an output display process 222 (e.g., for displaying on the ATM 204 to the customer). For system-level monitoring (e.g., system quality assurance) or other applications, the databases 210, 214, and 218 may be in communication with a master data server 224. In some implementations, the databases may be executed, for example, by a standalone computer system 226.

Processes in such multi-process processing systems may be executed in distinct execution phases. Execution of processes can include execution of one or more tasks within the processes in distinct execution phases. By segmenting execution into such phases, the distinct execution phases may be terminated, e.g., by multiple logical endpoints or breakpoints in data processing. Each execution phase may have one or more processing actions to achieve the objectives of that execution phase. As an example, the validate account process 208 may be executed in distinct execution phases in one or more manners. For example, as a first execution phase, the validate account process 208 may initially receive personal identification number (PIN) information from a customer. The various processing actions in receiving the PIN information from the customer can include, for example, displaying a prompt on the ATM display and running a routine to verify data entry of the PIN information. In a next execution phase, the process 208 may establish a connection to the database 210 and use the PIN information as a key to identify the customer's record. Once the process 208 has completed its transaction with

the customer's record, the connection to the database 210 may be terminated. In a final execution phase, the process 208 may generate results based on the foregoing transactions. As such, each of these execution phases include distinct logical endpoints (or processing breakpoint) at which the validate account process 208 may be temporarily suspended and/or resumed.

In some situations, one or more events may occur that tend to affect the course of normal system operation. Such events may be exceptions or errors raised by either hardware or software modules included in the processing system. For example, hardware exceptions or errors may include resets, interrupts or other signals from one or more hardware units. Exceptions may also be generated by an arithmetic logic unit for numerical errors such as divide by zero, overflow, instruction decoding errors, undefined instructions, etc. Based upon the occurrence of one or more such events, it may be necessary to temporarily halt operations of one or more of the databases for taking corrective action (e.g., maintenance, switching-over to secondary system, etc.).

Other events that may call for halting of operations and corrective action may include detecting the failure of one or more of the databases 210-224. Such failures can occur for a variety of reasons. For example, there may be an error in memory allocation or a conflict in writing to a memory space. There may also be an error in an underlying data operation such as when a process attempts to withdraw funds from a depleted account. In addition to events where there are temporary failures, there may also be events that are triggered by operator intervention. In some implementations, the operator may correct the situation that caused the failure or the system may, in time, correct the situation. Examples of events can include, without limitation, a failure of one or more devices connected to a network, a shut down of one or more devices or software services for maintenance, a failure and switch over of a device or software service, an exhaustion of resources such as storage space, an overload of processing units, a time-out of one or more software services.

To detect and address such events, data processing systems may use one or more techniques often referred to as checkpointing techniques to ensure minimum system downtime in the event of a failure or when the system is taken offline for maintenance or switch-over. A checkpointing technique generally involves storing details of a current

state of a process as a checkpoint record such that the process may use the stored information to be later restarted from that state. For example, the validate account process 204 may save its current state in a checkpoint record upon the completion of each execution phases (prior to starting execution of the next execution phase or other processing).

A checkpoint record may include various types of information such as process values, information about successfully processed records and other details relevant to the current execution phase of the process. For example, a checkpoint record can include information about a current position in a data queue (e.g., data queue 112 of FIG. 1) from which data is being processed. As such, after halting operations, processing may be resumed from this queue position. Along these lines, after recovery from a system failure, the process is able to restart from the stored intermediate checkpoint state rather than restart from an initial state.

As an example, if the PIN database 210 fails, the validate account process 208 may raise an exception to cause the entire processing system to terminate. Upon restart, the processes (or a portion of the processes) in the processing system may continue processing from their last checkpoint states. In this example, since the failure and restart occurs at a point in time after the customer provided his PIN information, the PIN information is restored to the process and does not need to be recollected from the customer. As such, there may be no need to once again prompt the customer to provide his PIN information.

In the event of a system failure, terminating and restarting a processing system from the most recently complete checkpoint can consume an undue amount of processing time and resources. After a processing system has terminated its activity in response to the exception condition, the processing system may need to be manually reinitialized by an information technology specialist experienced with such systems. This can result in significant system downtime. In some examples, a separate process may need to be designed to detect the system failure and inform a specialist. Examples of checkpointing systems are described in U.S. Pat. No. 6,584,581 entitled “Continuous Flow Checkpointing Data Processing”, U.S. Pat. No. 5,819,021, entitled “Overpartitioning system and method for increasing checkpoints in component-based parallel applications,”

and U.S. Pat. No. 5,712,971, entitled "Methods and Systems for Reconstructing the State of a Computation", of which the contents of each are incorporated herein in their entirety.

To improve efficiency and reduce processing resource consumption, processes within the processing system are executed from their last recorded checkpoints instead of restarting the entire system upon a failed connection to a process being restored. In one implementation, rather than terminating and restarting the entire system, the individual processes in the system may be informed to suspend processing until the failed connection is restored.

FIG. 3 shows a block diagram of a multi process system 300 that includes a data source process 302, processes 304a-n, a data sink process 306, and a fault-tolerance manager 308, which is in communication with each of the other processes (processes 302, 204a-n). In some implementations, the fault-tolerance manager 308 may be executed as another process within the multi process system 300. In some situations, the fault-tolerance manager 302 may be an application running on a separate computer system (not shown) or implemented in a dedicated processor, such as the checkpoint processor described in, for example, U.S. Pat. No. 6,584,581 entitled "Continuous Flow Checkpointing Data Processing", the contents of which are incorporated herein in their entirety.

One or more techniques may be implemented to establish communication among the processes 302-306 and the fault-tolerance manager 308. For example, individual exception channels 310a-n may be used for communicating information about exception conditions that may occur in the processes 302-306. The channels 310a-n can be part of a wired, wireless or combination wired and wireless network system. The channels 310a-n may be used by the processes 302-306 to communicate error information about the processes 302-306 to the fault-tolerance manager 308. For example, if an external device in communication with the process 304a should fail, the process 304a can immediately raise an error flag and communicate the error to the fault-tolerance manager 308 over the exception channel 310b.

In addition to exception channels 310a-n, the fault-tolerance manager 308 may send command messages (e.g., checkpoint command messages) to the processes 302-306 through the corresponding communication channels 312a-e. The communication

channels 312a-e are arranged to transmit command messages from the fault-tolerance manager 308 sequentially to each of the processes 302-306. For example, a message from the fault-tolerance manager 308 may be first communicated to the data source process 302 and then serially passed through each of the processes 304a-n and the data sink process 306 through the channels 312b-d. The data sink process 306 may use the channel 312e to communicate command messages to the fault-tolerance manager 308.

To store respective checkpoint data associated with each of the processes 302, 304a-n, 306, a storage area (e.g., memory) may be assigned to each process. Each process periodically suspends its current operation at the end of a distinct execution phase and stores its checkpoint data in the associated storage area. For example, the data source process 302 may periodically suspend its current operation at the end of distinct execution phases in processing (such as logical breakpoints in the stream of incoming data) and stores checkpoint information in the storage area 312. In this manner, as each of the processes 302, 304a-n, 306 are executed, corresponding storage areas 312, 314a-n, and 316 periodically saves checkpoint data. The storage areas 312-316 may be implemented with various types of storage techniques such as on non-volatile storage, e.g., magnetic media, such as a hard disk drive. In some examples, a single storage area may be shared by two or more processes. The checkpoint data may include information about current states and/or data associated with the processes 302-306 to allow for reconstruction of those states at a later time.

In some examples, the data storage areas 312-316 may also store the data queues associated with the processes (e.g., queues 110, 112, and 114 of FIG. 1). For example, the storage area 316 may include a queue for receiving processed data from process 304n and from which data can be output or published (e.g., printed, or displayed).

The fault-tolerance manager 308 manages the checkpointing operation by generating and sequentially passing checkpoint command messages through the communication channels 312a-e to each process 302-306. The checkpoint command messages pass through each process 302-306 so that the process may checkpoint its current state upon receiving the message. As such, the checkpoint command message travels to the data source process 302 and then sequentially passes through each process 304a-n and data sink process 306 before being returned to the fault-tolerance manager

308. This checkpointing operation may be automatically initiated at regular intervals. For example, the fault-tolerance manager 308 may initiate checkpoint command messages at a predetermined periodic rate, e.g., every five minutes. The periodic rate may be set to a default value, or adjusted by a user.

One or more techniques may be implemented for initiating operations for storing the checkpoint for processes included in system. For example, one or more external triggers may initiate operations for storing checkpoint information. In one instance, a network message may inform the fault-tolerant manager 308 of an impending network shutdown and thus trigger a checkpointing operation. In some implementations, the checkpointing operation may be triggered in response to values within or derived from the data records being processed. For example, the processed data records may include timestamps or breakpoint values that can be considered as logical points at which checkpointing may occur.

Along with storing checkpoint information during periods in which data is being processed by the system, information may be stored prior to processing data. In one implementation, an initial checkpointing operation may be triggered upon the multi-process system 300 being first initialized, e.g., during start-up. The fault-tolerance manager 308 may pass an initial checkpoint command message through each of the processes 302-306. In the example shown in FIG. 3, the initial checkpoint message is first communicated to the data source process 302. The data source process 302 immediately checkpoints, e.g., stores data that represents its initial state to the associated data storage space 312 and passes the initial checkpoint message downstream to the next process 304a. This initial checkpoint state is referred to as checkpoint state zero.

Similarly, in a serial manner, each of the processes 304-306 may correspondingly stores its initial state and associated data values to the appropriate storage area as checkpoint state zero. In examples, the initial state and associated data values can include initial values of global variables, reference data information, and auditing variables including initial values of counters.

After each of the processes 302-306 have stored their initial states, the initial checkpoint command message is returned to the fault-tolerance manager 308 through the channel 312e. Based upon the message being returned to the fault-tolerance manager 308

after its round trip through the processes 302, 304a-n, 306, the fault-tolerance manager is alerted that the processes 302-306 have completed checkpoint state zero. In some implementations, while the downstream processes are saving their current states, the source and other upstream processes may continue to receive data and perform other functions without waiting for all processes to save their states.

Similarly, additional checkpointing may be performed for each distinct execution phase of the processes 302-306. As such, along with storing data that represents the initial checkpoint information, the fault-tolerance manager 308 may initiate the storage of additional information, for example, which represents information associated with subsequent checkpoint cycles (e.g., checkpoint states 1, 2, 3, ... n). To initiate storage of subsequent checkpoint information, techniques such as propagating further checkpoint command messages through the processes 302, 304a-n, 306 may be utilized. Upon receipt of a checkpoint command message, the process 304a may either complete any ongoing tasks or suspend any outstanding tasks. In some examples, the process 304a may delete previously created checkpoint records stored in the data storage 314 and reclaim the storage space. The process 304 can then create a new checkpoint record for its current state and associated data. In some scenarios, earlier checkpoint records are persistently stored in memory and not overwritten by new checkpoint records. Further examples of information stored in checkpointing records are provided in U.S. Patent No. 6,584,581, the contents of which are incorporated herein in its entirety.

In some implementations, the process 304a may be in further communication with an external database 318 (executed on computer system 320). Sometimes, the connection to the database 318 may fail or the database 318 may be taken offline for maintenance. The failure could be a hardware failure of the computer system 320 executing the database 318. In such situations, the process 304a may raise an error flag over the exception channel 310a to notify the fault-tolerance manager 308 of a loss of connection.

Upon receiving notification of the error, the fault-tolerance manager 308 may generate and propagate an abort command message through the processes 302-306. The abort command message is typically first communicated to the data source process 302 by the channel 312a, then through each of the processes 302-306 by channels 312b-d, and finally back to the fault-tolerance manager 308 by channel 312e. Upon receiving the

abort command message, each of the processes 302-306 aborted its current activity with a relatively small delay (if any) and flushes/discards any outstanding tasks or records that may have been processed since the last checkpoint state. After a process has aborted activity, it may pass the abort command message to the next downstream process. In this manner, the abort command message propagates through to the sink process 306 before being returned to the fault-tolerance manager 308. The fault-tolerance manager 308 waits until it receives the abort command message from the sink process 306, which establishes that all of the processes 302-306 have aborted current processing tasks (e.g., are in a quiescent state).

In the scenario in which the database 318 has failed due to a hardware failure in the computer system 320, the processes 302-306 are directed to abort their processing. In some implementations, after the system has fully aborted its processing, the process 302 may wait for a specifiable amount of time, a time that should reflect the average amount of time that it takes to correct the failure, and once again begin processing from the last saved checkpoint state. In some implementations, the process 304a may periodically poll the database 318 to check for its status (i.e., to check whether the database 312 is operational). In some examples, the computer system 320 may be configured to automatically notify the process 304a when the database 318 is restored to operational state. When the connection with the database 318 is restored, the processing system 300 may once again begin processing from the last saved checkpoint state.

In this regard, the process 304a notifies the fault-tolerance manager 308 that the connection has been restored. The fault-tolerance manager 308 determines the last successfully completed checkpoint state for each of the processes 302-306, and sends a resume processing message to each of the processes 302-306. As with the other command messages, the resume processing message is propagated over the communication channels 312a-e sequentially through each of the processes 302-306. The resume processing message specifies the checkpoint state from which the processes 302-306 are to resume processing.

In some instants, the fault-tolerance manager 308 may initiate additional checkpointing operations while prior checkpointing operations are currently being executed. For example, while process 304n is processing an arbitrary checkpoint state

(e.g., checkpoint state N corresponding to checkpoint command message N), the fault-tolerance manager 308 may begin a subsequent checkpoint state (e.g., checkpoint state N+1) by generating and transmitting a subsequent checkpoint command message N+1 to the source process 302. Along these lines, it is possible that as checkpoint command message N is still traveling through the processes 302-306, a new subsequent checkpoint command message N+1 is generated and passed through the processes 302-306. In this manner, the fault-tolerance manager 308 can cause more frequent checkpointing of process states without having to wait until previous checkpointing states are completed.

In some situations, it is possible that a system failure may occur while one or more checkpoint command messages are in transit through the processes 302, 304a-n, 306. For example, consider a scenario in which the fault-tolerance manager 308 has initiated checkpoint state N by generating checkpoint command message N. While the checkpoint command message N is being processed by the processes 302-306, the connection between one of the processes (e.g., process 304a) and an external system (e.g., database 312) may fail. Upon being alerted to the situation, the fault-tolerance manager 308 may respond by passing an abort command message through the processes 302-306. The abort command message may reach a process (e.g., process 304n) that is still processing checkpoint state N (e.g., storing checkpoint information associated with the checkpoint N). Based upon receipt of the abort command, the process 304n may take one or more actions. For example, the process 304n may complete checkpoint state N and abort all further processing. In another scenario, the process 304n may immediately discard results associated with the current and subsequent states since a previous checkpoint state N-1 and abort further processing. As a result, when the system 300 achieves quiescence, each of the processes 302-306 may be at different checkpoint states. For example, all processes that are upstream from process 304n (e.g., the data sink process 306) may have completed checkpoint state N, while all processes downstream (e.g., process 304a and the data source process 302) from process 304n may have completed only checkpoint state N-1.

When the system 300 is ready to resume processing, the fault-tolerance manager 308 transmits one or more resume processing message to each process over the communication channels 312a-e. The resume processing message indicates to the

processes the earliest, fully committed (or completed) checkpointed state (e.g., checkpoint state N-1) from which they are to execute. In examples, processes that may have already completed checkpoint state N may simply reproduce the results from checkpoint state N-1 to checkpoint state N. In this manner, the processes 302-306 may avoid duplicating their earlier efforts. In examples, replaying the results from checkpoint state N-1 to checkpoint state N involves reproducing the results of the earlier processing actions that may have occurred between the two checkpoint states.

In examples, a system failure can occur substantially immediately after start-up. In such situations, many of the processes 302-306 may have only completed checkpoint state zero. These processes 302-306 can resume processing from checkpoint state zero based on initialization data and start-up values stored in the corresponding checkpoint records.

FIG. 4 is a flowchart depicting exemplary execution of a process (e.g., process 302 of FIG. 3) within a multi-process system. At start-up, the process immediately stores its initial state to data storage as checkpoint state zero (Step 402). The process may then be executed in distinct execution phases (e.g., execution phase 1, 2, ... N-1). At the conclusion of each execution phase, the process may save its end state to data storage, for example, the end state may be saved as a checkpoint state. For example, after a first execution phase, the process may save the end state of the first execution phase as checkpoint state 1 (Step 404). Similarly, after subsequent execution phases, the process may save the end states of the execution phases as checkpoint states 2, ... N-1, and N (Steps 406-410).

FIG. 5 is a flowchart depicting example steps executed in the event an external system fails or is taken offline for maintenance. For example, the external system may be a database (e.g., database 312 of FIG. 3) in communication with a process in a processing system. When the external system is taken offline for maintenance, or fails, an error flag may be raised by, for example, a process in communication with the failed external system (Step 502). An abort command message may be generated and propagated through the processes in response to the error flag (Step 504). The current activity of each of the processes is aborted when the process receives the abort command message (Step 506). Also, any transactions performed since the last checkpoint state may be

discarded by the processes. Further action is then aborted until the failed connection to the external system is restored (Step 508). When the connection is restored, a resume processing message is transmitted to each of the processes (Step 510). The resume processing message indicates the checkpoint state from where the processes are to resume processing. As such, each of the processes retrieves the relevant information regarding the checkpoint state from their associate storage areas (Step 512).

FIG. 6 is a flowchart depicting example steps executed in storing and resuming from checkpoints while executing a process. For example, upon initializing the process, information relating to an initial state of the process is stored in an associated storage area (Step 602). The process is then executed in distinct execution phases. As such, at the end of each execution phase, the process stores information representative of an end state of the execution phase (Step 604). When a predetermined event occurs, e.g., loss of connection to an external device (Step 608), execution of the process is aborted (Step 606). In the meantime, the process checks to see if the event that triggered the abortion has cleared (e.g., restoration of connection to the external device). During this time, the process is not shut down, but processing is aborted until the event is deemed cleared. Execution of the process is resumed from the last saved initial or end state (Step 610).

The techniques described above can be implemented using software for execution on a computer. For instance, the software forms procedures in one or more computer programs that execute on one or more programmed or programmable computer systems (which may be of various architectures such as distributed, client/server, or grid) each including at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. The software may form one or more modules of a larger program, for example, that provides other services related to the design and configuration of dataflow graphs. The nodes and elements of the graph can be implemented as data structures stored in a computer readable medium or other organized data conforming to a data model stored in a data repository.

The software may be provided on a storage medium, such as a CD-ROM, readable by a general or special purpose programmable computer or delivered (encoded in a propagated signal) over a communication medium of a network to the computer

where it is executed. All of the functions may be performed on a special purpose computer, or using special-purpose hardware, such as coprocessors. The software may be implemented in a distributed manner in which different parts of the computation specified by the software are performed by different computers. Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described.

It is to be understood that the foregoing description is intended to illustrate and not to limit the scope of the invention, which is defined by the scope of the appended claims. For example, a number of the function steps described above may be performed in a different order without substantially affecting overall processing. Other embodiments are within the scope of the following claims.

CLAIMS

1. A computer-implemented method, including:
storing information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase;
aborting execution of the process in response to a predetermined event; and resuming execution of the process from one of the saved initial and end states without needing to shut down the process.
2. The computer-implemented method of claim 1 wherein the predetermined event represents loss of connection to an external device.
3. The computer-implemented method of claim 1 wherein the predetermined event represents an error with an external device.
4. The computer-implemented method of claim 2 wherein the execution of the process is resumed when the connection to the external device has been restored.
5. The computer-implemented method of claim 3 wherein the execution of the process is resumed when the error with the external device has been cleared.
6. The computer-implemented method of claim 1 wherein the execution of the process is resumed from an end state that is stored prior to an execution phase in which the predetermined event occurred.
7. The computer-implemented method of claim 1 wherein the execution of the process is resumed from the initial state if the predetermined event occurs during an execution phase substantially immediately after startup of the process.

8. The computer-implemented method of claim 1 wherein execution of an execution phase includes performing one or more processing actions on a received stream of data to produce output data corresponding to the execution phase.

9. The computer-implemented method of claim 8, further including: storing the output data corresponding to one or more execution phases, and reproducing the output data when execution of the process is resumed.

10. The computer-implemented method of claim 1 wherein the process is part of a process stage and in communication through a data path to a second, different process in the process stage.

11. The computer-implemented method of claim 10, further including: passing a checkpoint command message through each process in the process stage; and

at each process, saving new information related to an initial or end state upon receiving the checkpoint command message, wherein the saving includes suspending operation of the process and saving the new information to a storage area.

12. The computer-implemented method of claim 11, further including: overwriting an old saved initial or end state with the new initial or end state.

13. The computer-implemented method of claim 11, wherein each process is in communication with a data queue for receiving and queuing data for the process.

14. The computer-implemented method of claim 11, further including: generating the checkpoint command message in response to detecting a trigger event.

15. The computer-implemented method of claim 14 wherein the trigger event includes information about a network event.

16. The computer-implemented method of claim 11, further including:
periodically generating the checkpoint command message.
17. The computer-implemented method of claim 11, further including:
generating the checkpoint command message in response to an occurrence of selected data values within or derived from incoming data records being processed.
18. The computer-implemented method of claim 11 further including:
generating an abort command message in response to the predetermined event;
passing the abort command message through each process in the process stage;
and
on receiving the abort command message at each process, aborting execution of the process and passing the abort command message to a next process.
19. The computer-implemented method of claim 1 further including:
resuming execution of the process from a selected one of the saved initial or end states based in part on information contained in a resume processing message.
20. The computer-implemented method of claim 17 further including:
receiving the abort command message during a first execution phase substantially immediately after the process is initialized and information related to the initial state of the process has been saved; and
resuming execution of the process from the saved initial state without need to shut down and restart the process.
21. A computer-readable storage medium storing a computer program including instructions for causing a computing system to:
store information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and

upon completion of the executing of the execution phase storing information representative of an end state of the execution phase;

abort execution of the process in response to a predetermined event; and
resume execution of the process from one of the saved initial and end states without needing to shut down the process.

22. A computing system, including:

an input device or port configured to receive and store information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase; and

at least one processor configured to:

abort execution of the process in response to a predetermined event, and
resume execution of the process from one of the saved initial and end states without needing to shut down the process.

23. A computing system, including:

means for storing information related to an initial state of a process upon being initialized, wherein execution of the process includes executing at least one execution phase and upon completion of the executing of the execution phase storing information representative of an end state of the execution phase; and

means for controlling execution of the process, the controlling including:

aborting of execution of the process in response to a predetermined event, and

resuming execution of the process from one of the saved initial and end states without needing to shut down the process.

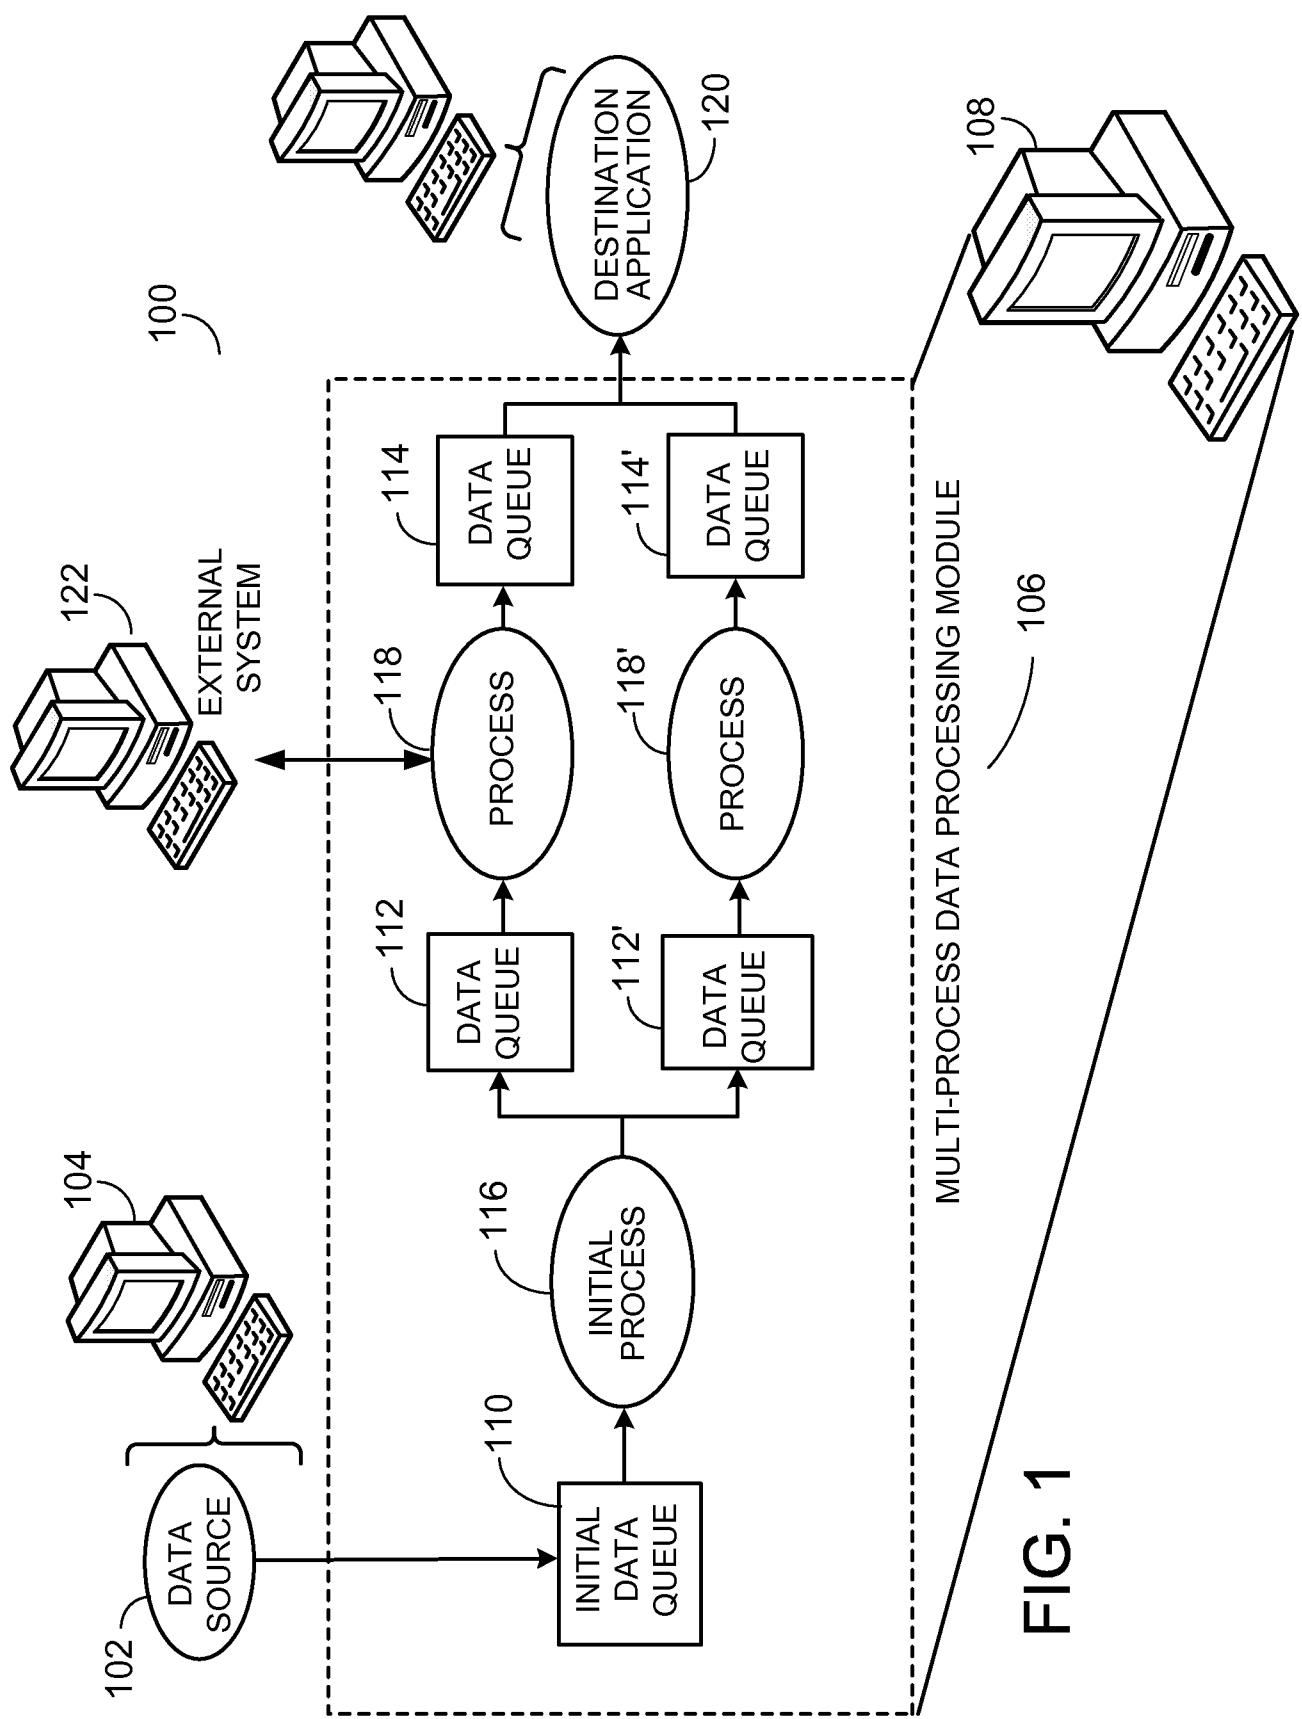


FIG. 1

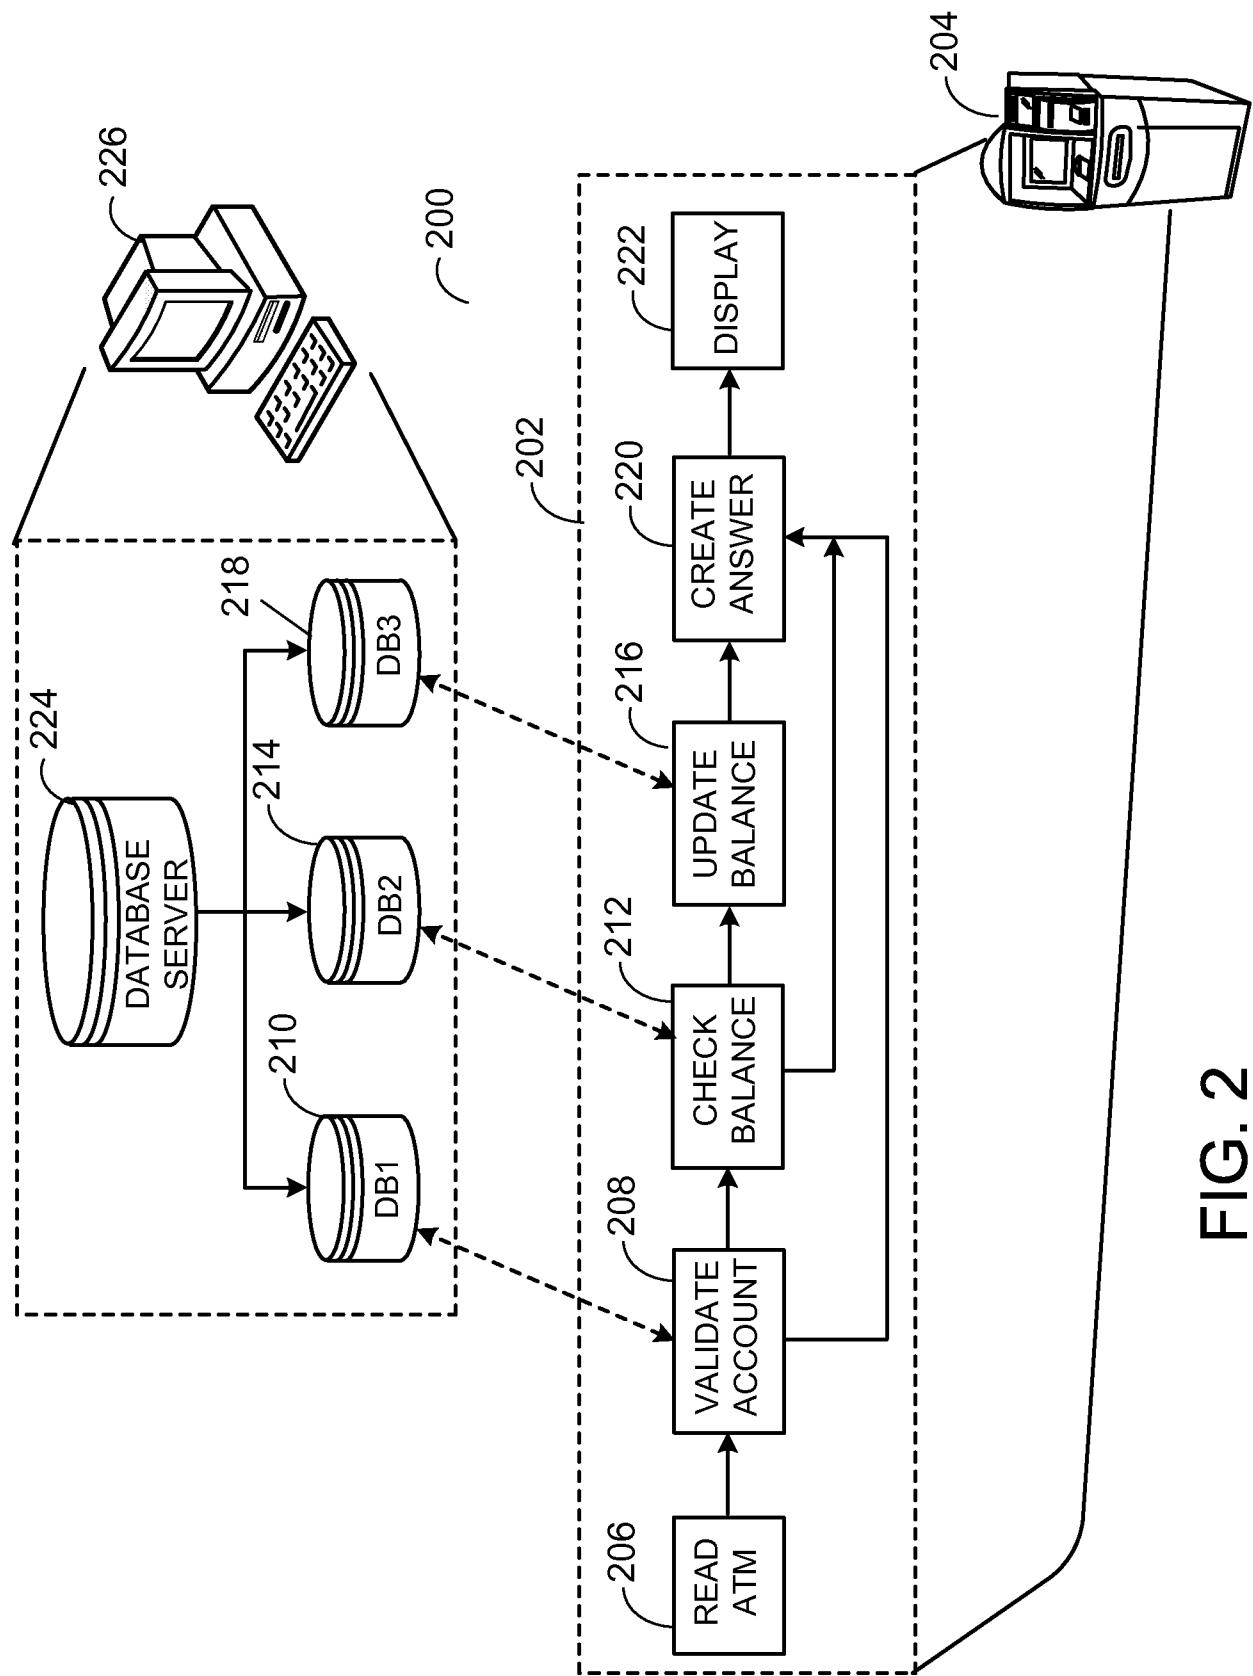


FIG. 2

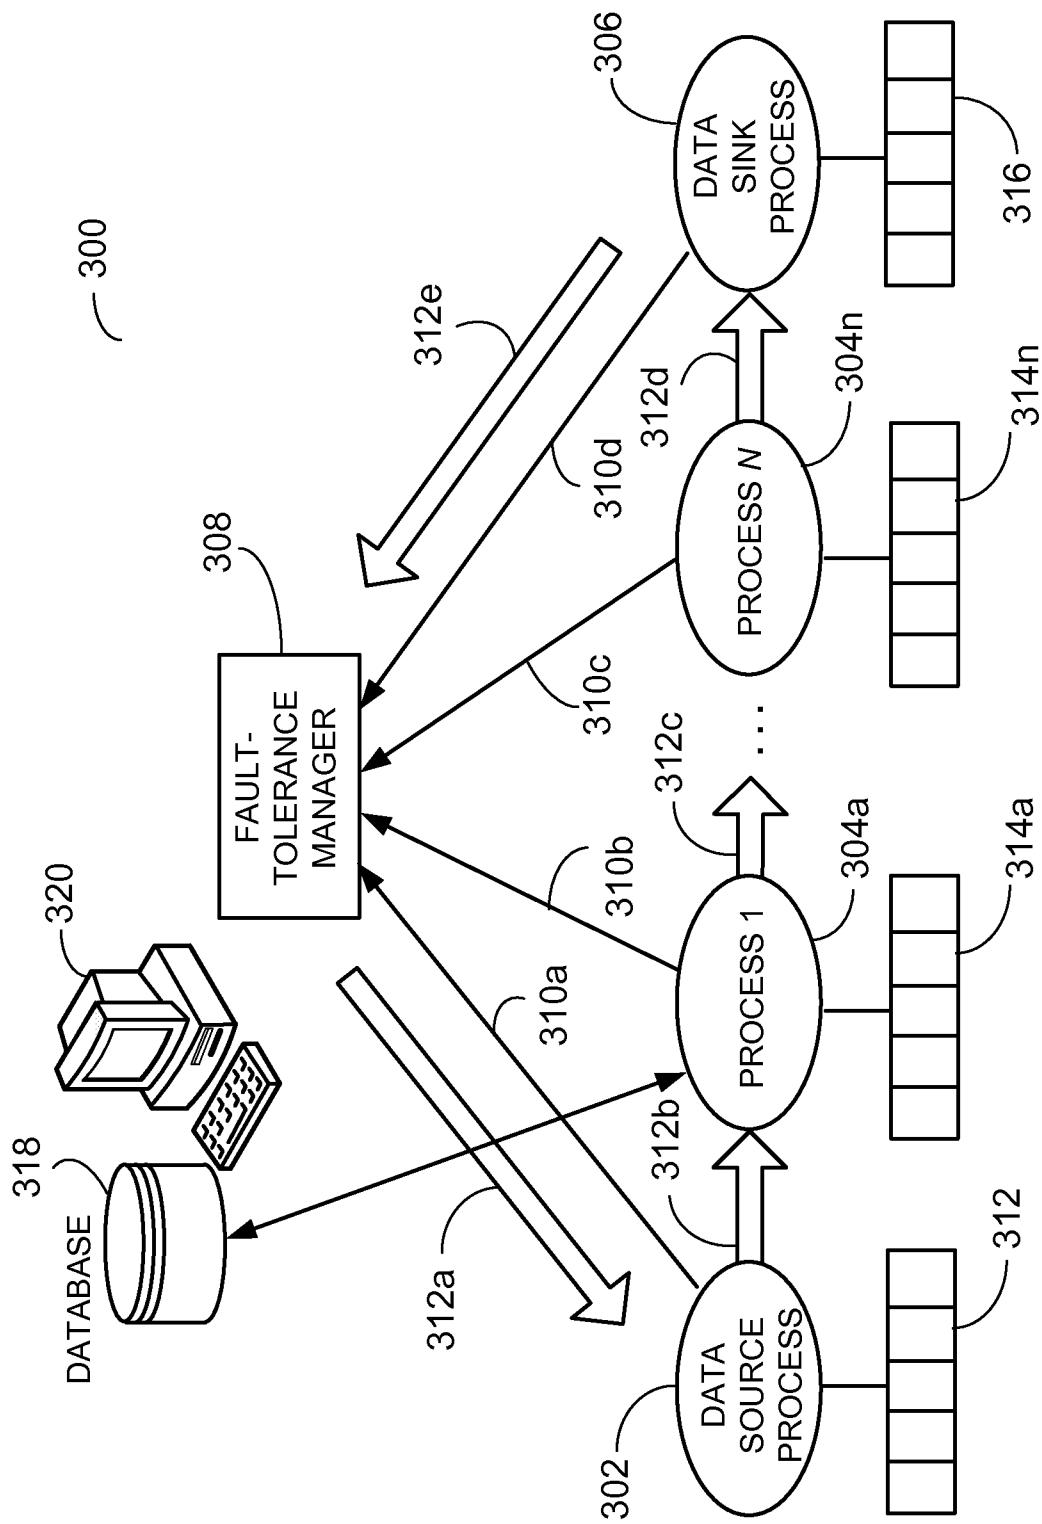


FIG. 3

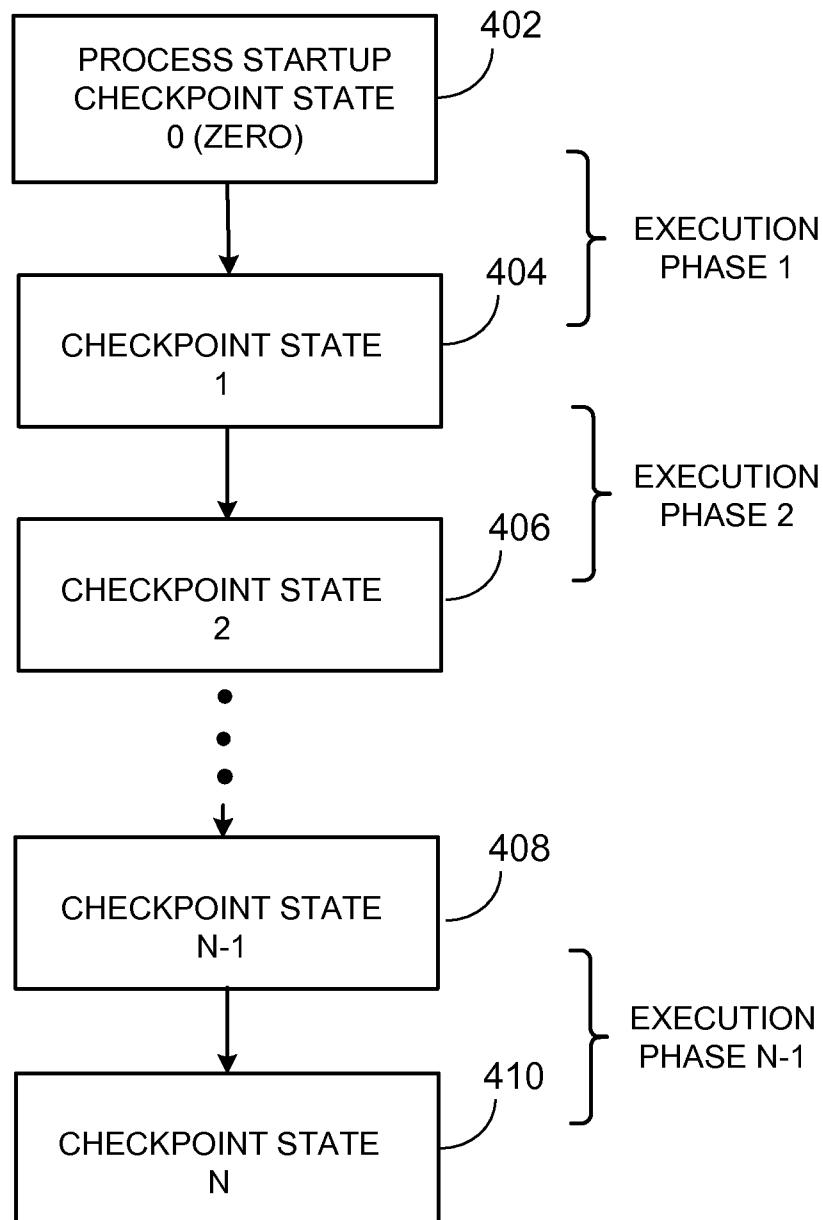


FIG. 4

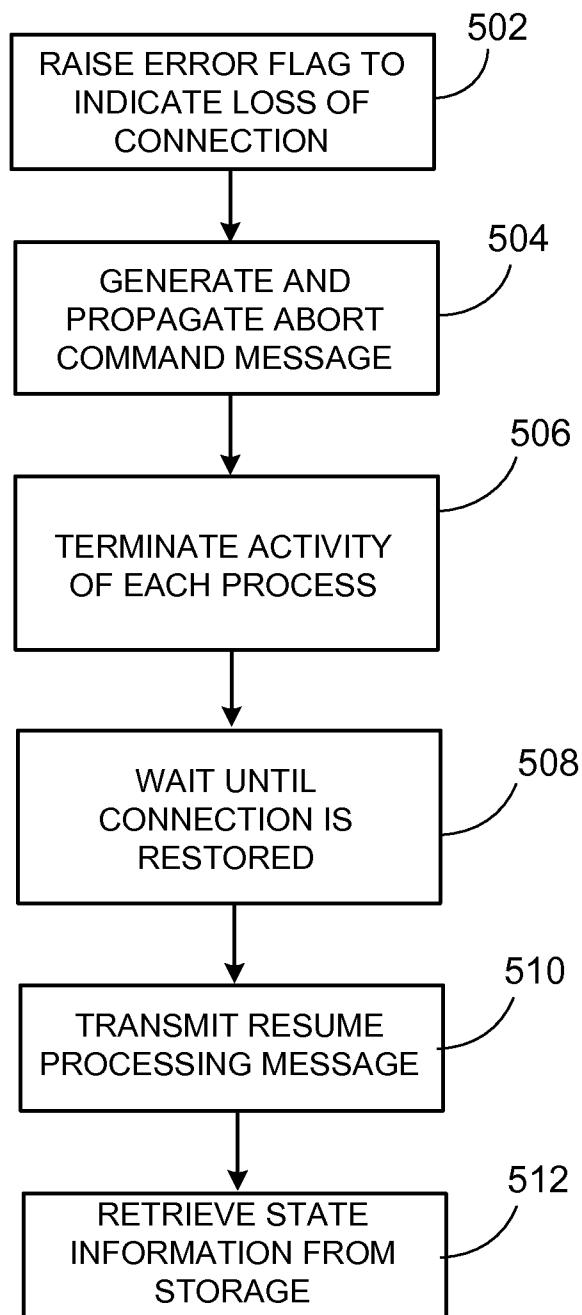


FIG. 5

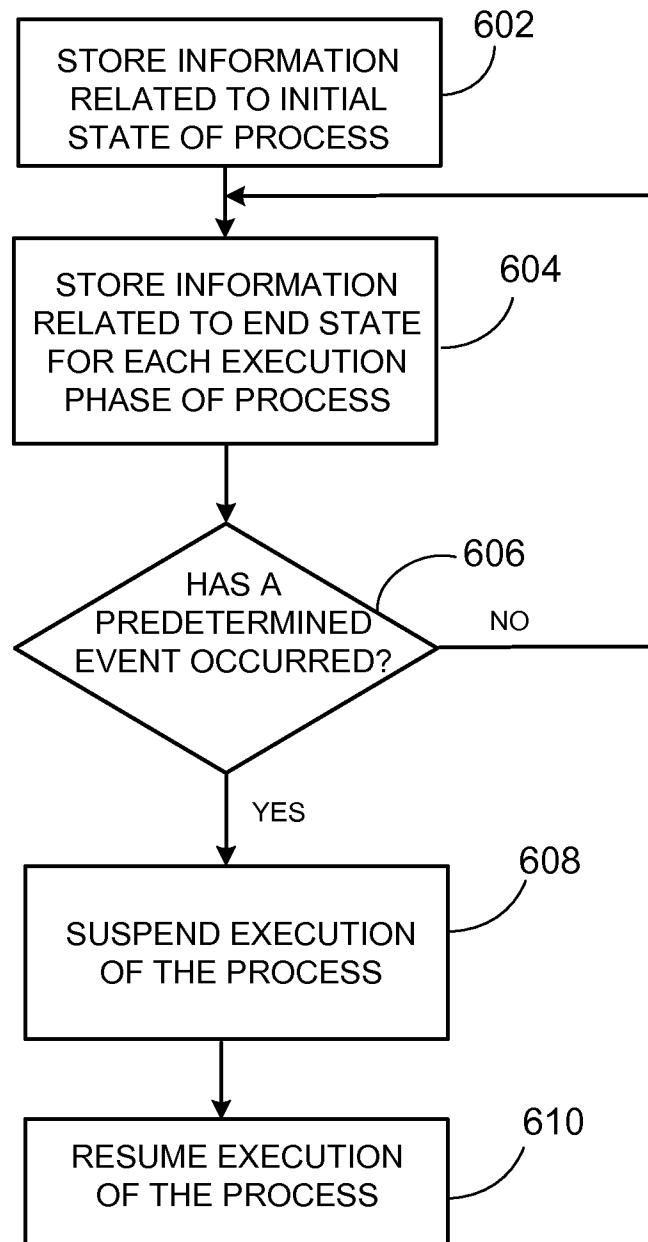


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/025388

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F11/14
ADD. G06F9/48

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | <p>R. BALDONI ET AL: "Characterization of consistent global checkpoints in large-scale distributed systems", PROCEEDINGS OF THE FIFTH IEEE COMPUTER SOCIETY WORKSHOP ON FUTURE TRENDS OF DISTRIBUTED COMPUTING SYSTEMS, 1 January 1995 (1995-01-01), pages 314-323, XP55026754, DOI: 10.1109/FTDCS.1995.525000 ISBN: 978-0-81-867125-8 page 314, right-hand column, line 21 - page 318, right-hand column, line 22; figures 1, 2</p> <p style="text-align: center;">-----</p> <p style="text-align: center;">-/-</p> | 1-23 |

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

10 May 2012

Date of mailing of the international search report

22/05/2012

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

Michel, Thierry

INTERNATIONAL SEARCH REPORT

| |
|---|
| International application No PCT/US2012/025388 |
|---|

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|--------------------------|
| A | <p>HIGAKI H ET AL: "Checkpoint and rollback in asynchronous distributed systems", INFOCOM '97. SIXTEENTH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES. DRIVING THE INFORMATION REVOLUTION., PROCEEDINGS IEEE KOBE, JAPAN 7-11 APRIL 1997, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, vol. 3, 7 April 1997 (1997-04-07), pages 998-1005, XP010251940, DOI: 10.1109/INFCOM.1997.631114 ISBN: 978-0-8186-7780-9 page 999, right-hand column, line 1 - page 1002, left-hand column, line 25</p> <p>-----</p> <p>YUQUN CHEN ET AL: "CLIP: A Checkpointing Tool for Message Passing Parallel Programs", PROCEEDINGS OF THE 2ND ACM WORKSHOP ON ROLE-BASED ACCESS CONTROL. RBAC '97. FAIRFAX, VA, NOV. 6 - 7, 1997; [ACM ROLE-BASED ACCESS CONTROL WORKSHOP], NEW YORK, NY : ACM, US, 15 November 1997 (1997-11-15), pages 33-33, XP010893034, ISBN: 978-0-89791-985-2 page 35, left-hand column, line 5 - page 39, right-hand column, line 24</p> <p>-----</p> <p>Ferda Tartanoglu: "Using the B Method for the Formalization of Coordinated Atomic Actions", CS-TR-865, Technical Report Series, 1 October 2004 (2004-10-01), pages 1-18, XP55026374, University of Newcastle upon Tyne, UK Retrieved from the Internet: URL:http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.65.8623&rep=rep1&type=pdf [retrieved on 2012-05-07] page 8, line 1 - page 10, line 17</p> <p>-----</p> <p>US 5 819 021 A (STANFILL CRAIG [US] ET AL) 6 October 1998 (1998-10-06) cited in the application the whole document</p> <p>-----</p> | 1-23 |
| A | | 1,11, 14-18, 21-23 |
| A | | 1,10,11, 18,21-23 |
| A | | 1-23 |

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/025388

| Patent document cited in search report | Publication date | Patent family member(s) | | Publication date |
|--|------------------|-------------------------|-------------|------------------|
| US 5819021 | A 06-10-1998 | AT | 359551 T | 15-05-2007 |
| | | AU | 1286197 A | 17-07-1997 |
| | | CA | 2240286 A1 | 03-07-1997 |
| | | DE | 69637020 T2 | 16-08-2007 |
| | | DK | 0954781 T3 | 06-08-2007 |
| | | EP | 0954781 A1 | 10-11-1999 |
| | | ES | 2283008 T3 | 16-10-2007 |
| | | JP | 3207438 B2 | 10-09-2001 |
| | | JP | H11514116 A | 30-11-1999 |
| | | PT | 954781 E | 31-05-2007 |
| | | US | 5819021 A | 06-10-1998 |
| | | WO | 9723826 A1 | 03-07-1997 |