

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 February 2010 (25.02.2010)

(10) International Publication Number
WO 2010/022246 A2

- (51) **International Patent Classification:**
G06F 9/445 (2006.01) G06F 15/16 (2006.01)
G06F 9/06 (2006.01)
- (21) **International Application Number:**
PCT/US2009/054483
- (22) **International Filing Date:**
20 August 2009 (20.08.2009)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/090,534 20 August 2008 (20.08.2008) US
- (71) **Applicant (for all designated States except US):** WAL-
MART STORES, INC. [US/US]; 702 S.W. 8th Street,
Bentonville, AR 72716-8095 (US).
- (72) **Inventor; and**
- (75) **Inventor/Applicant (for US only):** VANDERPOOL,
Trey [US/US]; 5102 SW Blairemont Road, Bentonville,
AR 72712 (US).
- (74) **Agents:** BREWSTER, Alicia, L., W. et al.; Troutman
Sanders LLP, 600 Peachtree Street NE, Suite 5200, At-
lanta, GA 30308 (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, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ,
TM), European (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, 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:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted
a patent (Rule 4.17(ii))

Published:

- without international search report and to be republished
upon receipt of that report (Rule 48.2(g))

(54) **Title:** PROCESS AUTO-RESTART SYSTEMS AND METHODS




Fig. 1

(57) **Abstract:** Systems and methods for auto-restarting abnormally terminated processes are disclosed. An auto-restart system can include a parent task control block, a child process, and a shared resource. The parent task control block can spawn the child process. The child process can operate on the shared resource. When the child process finds the shared resource locked, the child process can terminate abnormally. The parent task control block can recognize the abnormal termination of the child process, and can automatically rollback and restart the child process. Accordingly, the child process can be restarted to operate on the shared resource without human intervention.

WO 2010/022246 A2

PROCESS AUTO-RESTART SYSTEMS AND METHODS

CROSS-REFERENCE TO RELATED APPLICATION & PRIORITY CLAIM

This application claims the benefit, under 35 U.S.C. § 119(e), of U.S. Provisional Application Serial No. 61/090,534, filed 20 August 2008, the entire contents and substance of
5 which is hereby incorporated by reference.

TECHNICAL FIELD

Various embodiments of the present invention relate generally to computer process management, and more particularly, to systems and methods for auto-restarting processes in a database environment or other environment having shared resources.

10

BACKGROUND

In a database system, multiple jobs, or processes typically run simultaneously. A database likely contains many records, and processes operating on the database may require access to the records. There may also be instances in which multiple processes require access to a single record. Some forms of access to records, such as basic read operations, do not require
15 an acting process to lock the record on which it is operating. Other forms of access, however, such as typical write operations, may require a process to exclude other processes from performing operations on the accessed record. The database system may implement a locking mechanism to enable processes to lock a shared resource while certain operations are performed on the locked resource.

20

While locking can be beneficial in computing environments where multiple processes run simultaneously on a set of resources, locking may sometimes lead to inefficiencies. For example, a first process may be required to wait for a second process to complete a task on a locked shared resource. While waiting, the first process may repeatedly check the availability of the shared resource until the shared resource becomes available again. As a result, processor
25 power is used inefficiently to ensure that the second process can proceed as soon as possible.

A database system may be constructed such that, if a process waits too long for a particular resource, that process is automatically terminated. This termination can conserve computing power by reducing the repetitive checking that occurs while the process waits for the

shared resource. As an unfortunate result, when the desired resource is available, the ceased process is no longer running, although it could conceivably access the resource. In conventional database systems, a database administrator may be required to manually restart the terminated process, so that the process can continue its tasks.

5 Manually restarting processes, however, can lead to some undesirable results. For example, manual restart allows for human error in restarting a process. Additionally, manual restart will likely not occur immediately after the desired resource becomes available and, as a result, the system is not as efficient as desired.

10 Accordingly, there is a need for systems and methods for sensing when a process has ended, and for automatically restarting such process. It is such systems and methods to which various embodiments of the present invention are directed.

SUMMARY

15 Briefly described, various embodiments of the present invention generally comprise auto-restart systems, methods, and utilities for detecting abnormal termination of processes, and restarting terminated processes. Embodiments of the present invention can be implemented in database systems or many computing environments where multiple processes run concurrently while operating on one or more shared resources.

20 According to an exemplary embodiment of the present invention, an auto-restart system can comprise one or more shared resources, a plurality of task control blocks, and a plurality of child processes.

The shared resources can be database records or other forms of data in a computing system. The shared resources are shared among the child processes, in that multiple child processes can concurrently attempt to access the shared resources.

25 Each task control block, or parent process, can perform a process call to start a child process. After calling the child process, the task control block can remain in a hung state until the child process terminates. When the child process terminates, the task control block can receive an exit code output by the child process when the child terminated. As a result, the task control block can detect when its child process terminates, and with what exit code the child process terminates. The task control block can determine whether, when, and under what

conditions to restart the child process. Without human intervention, the task control block can restart a terminated child processes automatically.

The child processes can be spawned by the task control blocks, and can be configured to operate on the shared resources. In some instances, a first child process will lock a shared
5 resource while performing an operation on the resource. In that case, if a second child process attempts to access the same shared resource in a manner that is prevented by the lock, the second child process can be unsuccessful in accessing the shared resource. The second child process may be terminated after a predetermined period of time, or after some other condition is met, while the second process unsuccessfully attempts to access the shared resource.

10 The task control block that spawned the second child process can sense termination of the second child process and, if appropriate, can restart the second child process.

These and other objects, features, and advantages of the auto-restart systems and methods will become more apparent upon reading the following specification in conjunction with the accompanying drawing figures.

15 BRIEF DESCRIPTION OF THE FIGURES

Fig. 1 illustrates a block diagram of an auto-restart system, according to an exemplary embodiment of the present invention.

Fig. 2 illustrates a conventional database management system.

20 Fig. 3 illustrates a flow diagram of an implementation of the auto-restart system, according to an exemplary embodiment of the present invention.

Fig. 4 illustrates a flow diagram of an auto-restart method, according to an exemplary embodiment of the present invention.

Fig. 5 illustrates an implementation of a call attach facility utilized in the auto-restart system, according to an exemplary embodiment of the present invention.

25 Fig. 6 illustrates a conventional use of an IBM DSN system, which can be adapted for use in the auto-restart system, according to an exemplary embodiment of the present invention.

DETAILED DESCRIPTION

Various embodiments of the invention are auto-restart systems and methods. To facilitate an understanding of the principles and features of the auto-restart systems and methods,

various illustrative embodiments are explained below. In particular, the invention is described in the context of being an auto-restart utility for restarting database processes that terminate abnormally. Embodiments of the invention, however, are not limited to database contexts, but can be used in many systems providing shared resources or a plurality of processes.

5 The components described hereinafter as making up various elements of the auto-restart systems and methods are intended to be illustrative and not restrictive. Many suitable components that would perform the same or similar functions as the components described herein are intended to be embraced within the scope of the invention. Such other components not described herein can include, but are not limited to, for example, components developed after
10 development of the invention.

Referring now to the figures, wherein like reference numerals represent like parts throughout the views, embodiments of the auto-restart systems and methods will be described in detail.

Fig. 1 illustrates a block diagram of an auto-restart system according to an exemplary
15 embodiment of the present invention. As shown in Fig. 1, the auto-restart system can comprise a processor 50, one or more shared resources 110, a plurality of task control blocks 120, and a plurality of child processes 130. These elements of the auto-restart system 100 will be described further below.

The below discussion, while provided in various sections, should be read as a whole, and
20 as applying to this entire disclosure and the various embodiments of the auto-restart system 100. Thus, discussion of one or more features and embodiments in a particular section can also be pertinent to other features and embodiments discussed in one or more other sections.

Database Management Systems

Various aspects of the auto-restart system 100 can be elements of, or can interact with, a
25 database management system 200. Fig. 2 illustrates a conventional database management system, which can provide a basis for the auto-restart system 100. As shown in Fig. 2, the database management system 200 can comprise at least one database server 210, a database 220 (see Fig. 1), a database administrator 230, and at least one application server 240. The database management system 200 can be accessible by remote computing devices 250 via a network 270.

30 The database server 210 can maintain the database 220 on one or more storage devices. The storage devices can be of various types, including, without limitation, magnetic, optical, and

magneto-optical storage devices, or flash memory. The database 220 can include a plurality of database records, which can act as the shared resources 110 in the auto-restart system 100.

The database administrator 230 can be responsible for operation and maintenance of the database server 210, the database 220, and the application server 240.

5 In a typical scenario, a database client, such as the remote computing device 250, can allow a user to enter data through a user interface, such as a web browser. This data can be sent to the application server 240 across the network 270. Base on the data received from the database client 250, the application server 230 can form a database query, such as a SQL query. The application server 230 can then forward the query to the database server 210, directly or
10 across the network 270. In some database management systems 200, the application server 240 can be integrated into the database server 210. After receiving the query, the database server 210 can validate the query and, if valid, perform the operations required by the query.

The database server 210 can normally, but need not, be programmed to accommodate multiple database clients 250, which can concurrently connect to the database server 210 with
15 different requests. These requests can include, for example, routine online transaction processing (“OLTP”) requests, in which relatively few database records need to be processed with sub-second response times. In some instances, the database clients 250 can also issue processor-intensive decision support system (“DSS”) requests, which may require hours to complete.

20 To support queries, the database server 210 can implement a set of database server processes. These can include, without limitation, various transaction management, scheduling and data management functions, query optimization, scan processing, join processing, aggregation processing, sort processing, convergence processing, final result set processing, logging, recovery, index management, concurrency control, buffer pool management, and
25 parallel query processing. To efficiently perform these or other processes of the database management system 100, the database server 210 can often run multiple processes simultaneously. Some or all of the running processes can be required to operate on the database records 110 concurrently. Because database records are shared between the processes, there can be instances in which multiple processes seek simultaneous access to a single database record
30 110.

Database systems 200 conventionally implement one or more methods for avoiding data corruption that can occur as a result of maintaining shared resources 110, such as database records. A conventional database can implement locking, in which a process can lock a shared resource 110 when performing certain types of data accesses, such as write operations. When a first process locks a shared resource 110, such as a database record, a second process can be unable to access the shared resource 110 for the purpose of certain operations, such as additional write operations. As a result, the database management system 200 can avoid corruption that can occur when multiple processes attempt to write to the same shared resource 110.

A database process encounters a database contention error when that process cannot proceed because a shared resource 110 that it must access is locked. A negative result of locking is that, when a database contention error occurs, a process may be required to wait for its turn to access the shared resource 110. While waiting, the process can repeatedly check whether the shared resource 110 is available. This repeated checking can inefficiently utilize other system resources, such as processor power.

Some conventional database systems 200 seek to conserve processor power by terminating database processes that cannot proceed because of a database contention error. Once terminated, a waiting process can no longer waste process power, but the process is also unable to complete its assigned tasks. In a conventional database management system 200, a database administrator 230 must manually select which processes to restart and, if applicable, must rollback each of such processes to a checkpoint from which they should proceed. This manual intervention can lead to undesirable manual errors. Additionally, because a process remains terminated until a human intervenes, a process that might otherwise be able to proceed cannot do so until manually restarted. As a result, such a process is delayed until human intervention occurs.

Exemplary Aspects of Auto-Restart Systems and Methods

Various embodiments of the present invention can reduce wasteful processor usage, while enabling processes to proceed in their tasks without human intervention. Referring back to Fig. 1, as mentioned above, an exemplary embodiment of the auto-restart system 100 can comprise a processor 50, one or more shared resources 110, a plurality of task control blocks 120, and a plurality of child processes 130.

Operations of the auto-restart system 100 and its components can be performed by, or in, a computer system. Accordingly, such operations can be stored on a computer-readable medium of expression and performed by a computer processor, such as the processor depicted in Fig. 1. Further, various modules and components associated with the auto-restart system 100 can be implemented as hardware, software, or a combination thereof.

The one or more shared resources 110 can be records in the database 220 or other forms of data in a computing system. The shared resources 110 can be shared among the child processes 130, such that multiple child processes 130 can simultaneously attempt to access the shared resources 110.

The processor 50 can initialize the plurality of task control blocks 120, and can execute one or more processes for operation of the task control blocks 120.

Each task control block 120, or parent process, can perform a process call to initialize a child process 130 for running a program of the database management system 100. After initiating a child process, the task control block 120 can remain in a hung state until its child process terminates. Alternatively, in a multi-threaded implementation of the auto-restart system, the task control block 120 can spawn multiple child processes 130, and can monitor each of its child processes 130 to detect termination. Regardless of whether a multi-threaded implementation is used, when the child process 130 terminates, the task control block 120 can receive an exit code output by the child process 130 upon termination. Accordingly, the task control block 120 can detect this termination.

The child processes 130 can represent various programs operating within the database management system 100. Each of the plurality of child processes 130 can be spawned by a task control block 120. When a task control block 120 calls a program, the called program can become a child process 130 of the auto-restart system 100. The child process 130 can be configured to execute a sequence of tasks of the called program, and some of these tasks may require access to the shared resources 110. While each child process 130 executes its assigned tasks, a checkpoint utility can create one or more checkpoints of the child process 130. The checkpoints can be snapshots of the current state of the child process 130, and can include SQL codes and current values of variables. The checkpoints can be used to return the child process 130 to a point before abnormal termination of the child process 130.

In some instances, a first child process 130 can lock a shared resource 110 before performing an operation on the resource 110. In that case, if a second child process 130 attempts to access the same shared resource 110 in a manner that is prevented by the lock, the second child process 130 can encounter a database contention error. If the database contention error remains unresolved after a predetermined number, or timeframe, of unsuccessful attempts to access the shared resource 110, the second child process 130 can be terminated. In that case, the child process 130 can terminate itself or, alternatively, the task control block 120 running the child process 130 can terminate the child process 130.

When a child process 130 terminates, the child process 130 can output an exit code. The corresponding task control block 120 can receive the exit code output by the child process 130 upon termination. As a result, a task control block 120 can detect when a child process 130 terminates, and with what exit code the child process 130 terminates. The task control block 120 can also detect the SQL codes that were being run by the child process 130 before the child process 130 terminated. Based, at least partially, on the exit code and the last SQL code of the terminated child process 130, the task control block 120 can determine whether, when, and how to restart the child process 130. For example, if a child process 130 exits with a GOBACK code, the task control block 130 can rollback the child process 130 before restarting the child process 130. Additionally, the last SQL code of a child process 130 can be further utilized by the task control block 130 to determine how long to wait before restarting the child process 130.

In some exemplary embodiments, the auto-restart system 100 can maintain a participation list or an exclusion list, which can be based on the last SQL code of the child processes. The task control block 120 can compare a child process's last SQL code with the participation, or exclusion list. If a participation list is used, the auto-restart system 100 can restart the child process 130 if the last SQL code is on the participation list, and can select not to restart the child process if the last SQL code is absent from the participation list. Alternatively, if an exclusion list is used, the task control block 120 can restart the child process unless the last SQL code of the child process appears on the exclusion list

Fig. 3 illustrates a flow diagram of an implementation of the auto-restart system 100, according to an exemplary embodiment of the present invention. As shown in Fig. 3, an implementation of the auto-restart system can comprise various components of the task control block 120 and the child process 130. For example, and not limitation, the task control block 120

can include WMDB2ARF 310 and DSNALI 320, and the child process 130 can include WMDB2AR1 330, a COBOL program 340, DSNHLI 350, and DSNHLI2 360. WMDB2ARF represents an auto-restart facility 310, which can provide some or all of the primary functionality for the auto-restart system 100.

5 An exemplary embodiment of the auto-restart system 100 can utilize DB2 370, an IBM-registered database management system 200, which can run on the z/OS platform. The auto-restart system 100, however, is not limited to the DB2 370 database management system 200, and other database management systems 200 can be used with appropriate implementation modifications. DSNALI 320 and DSNHLI2 360 are two modules supplied by IBM as part of a
10 call attach facility (“CAF”) 500 (see Fig. 5). DSNALI 320 can be the main application program interface (“API”) used to connect to the DB2 370 database management system 200. DSNHLI2 360 can be the API used to make SQL calls to the DB2 370 database management system 200.

Fig. 5 illustrates a call attach facility 500, which can be utilized in an exemplary embodiment of the auto-restart system, 100. In a standard call attach facility 500
15 implementation, the child process 130 can make a direct call to DSNHLI2 360 to issue a DB2 370 database request, such as a SELECT, INSERT, UPDATE or DELETE request.

In a non-call attach facility environment using a DSN connection, the programs 340 executed by child processes 130 can be compiled to make direct calls to DSNHLI 350 to issue the DB2 370 database requests. The auto-restart system 100 can eliminate a need for job control
20 language (“JCL”) and child process 130 changes. Thus, the child processes 130 are compiled to make calls to the DSNHLI 350 API. The DSNHLI 350 can be a standard protocol to issue requests in the database management system 200 via the DSN system 610 provided by IBM.

Fig. 6 illustrates a conventional use of the DSN system 610, which can be adapted for use in an exemplary embodiment of the auto-restart system 100. The DSN system 610 can use
25 IBM’s DSNELI 620 main entry point to connect to the DB2 370 database management system 200. For more information refer to IBM’s documentation on DSNELI 620, DSNALI 320, DSNHLI 350, and DSNHLI2 360.

Referring back to Fig. 3, in an exemplary embodiment of the auto-restart system 100, DSNHLI 350 can be the main entry point for the DB2 370 database management system 200
30 requests from the child processes 130, including compiled child processes 130 and system utilities. The DSNHLI 350 module need not be loaded into memory. Instead, DSNHLI 350 can

be masked by a sub-routine WMDB2AR9 355 as part of the auto restart system 100. During initialization of the auto restart system 100, the load module WMDB2AR9 355 can be loaded into memory, and can be renamed to DSNHLI. The rename of WMDB2AR9 355 to DSNHLI need be only in memory in the current task control block 120 environment. As part of
5 initialization, WMDB2AR9 355 can receive the address of DSNHLI2 360. As renamed to DSNHLI, WMDB2AR9 355 can become a pass-through mechanism called for DSNHLI 350 requests. WMDB2AR9 355 can pass control to DSNHLI2 360 to support the standard call attach facility 500 protocol.

WMDB2AR9 355 can receive a request from a child process 130 via a DSNHLI 350 call,
10 and can pass the request forward to DSNHLI2 360 (IBM's call attach facility 500) to perform the DB2 370 database management system 200 operation. Upon return from DSNHLI2 360, WMDB2AR9 355 can save a SQLCA (IBM's SQL Communication Area) (not shown in the figures) control block in the auto restart system 100 memory, and can then return control back to the child process 130. The SQLCA can contain the SQL code as returned from the DB2 370
15 database management system 200. By saving the SQLCA in the auto restart system 100 memory, the SQLCA can be used to evaluate the last SQL code from the DB2 370 database management system 200 when a child process 130 terminates with a non-zero value exit code.

WMDB2AR1 330 can open a connection to DB2 370, load the DSNHLI2 360 load module, load the WMDB2AR9 355 load module, rename WMDB2AR9 355 to DSNHLI 350,
20 and call the requested child process 130 program. In other words, WMDB2AR1 330 can set up the DB2 370 database management system 200 environment for the child process 130 to communicate with the DB2 370 database management system 200.

In some exemplary embodiments, the auto restart system 100 can automatically terminate a child process 130 when the auto-restart system 100 detects a SQL code returned from
25 DSNHLI2 360 that appears in the participation list. This feature varies from previously-described operation of WMDB2AR9 355, which can generally return control back to the child process 130 for evaluation of the SQL code.

Also shown in Fig. 3, the Auto Trap option 380 as can provide a means to allow the auto-restart system 100 to detect the SQL codes, check the SQL code participation list, and then
30 terminate the child process 130 if the SQL code appears in the participation list. When the Auto Trap feature 380 is active and an SQL code in the participation list is detected, the auto-restart

system 100, specifically WMDB2AR9, can terminate the child process 130 with a user abend of U0102. In that case, control need not be returned to the child process 130. The user abend can terminate the child process 130, and thus, the task control block 120 can detect the termination and issue a restart on the child process 130.

5 Fig. 4 illustrates a flow diagram of a method 400 of the task control block 120 for auto-restarting a child process 130 in an exemplary embodiment of the auto-restart system 100. As shown in Fig. 4, the task control block 120 can spawn a child process 130 at 410. The child process 130 can provide various functionalities of the database management system 200. At 420, the task control block 120 can monitor the spawned child process 130. If the child process 130
10 terminates abnormally, such as by a database contention error, the child process 130 can output an exit code indicating the abnormal termination. The task control block 120 can receive the exit code and detect termination of the child process 130 at 430. At 440, the task control block 120 can analyze the exit code and the SQL code of the child process 130 to determine whether to restart the child process 130. If it is determined that the child process 130 should be restarted, at
15 450, the task control block 120 can rollback the child process 130 to a checkpoint. At 460, the task control block 120 can automatically restart the child process 130. Accordingly, the child process 130 can continue executing the tasks assigned to it for operation of the database management system 200.

Exemplary Implementation

20 Embodiments of the auto-restart system 100 can be implemented in various manners. For example, and not limitation, aspects of the auto-restart system 100 can be implemented through one or more COBOL programs. In that case, primary functionality of the auto-restart system 100 can be controlled via global and local parameter files, which are also known as parm files. Additionally, COBOL programs, or other types of program for operations of the database
25 management system 100 can be child processes 130 initiated by the task control block 120.

 The auto-restart system 100 can display error information when child processes 130 terminate, and can issue rollbacks through an exit sub-routine based on checkpoints. The checkpoints can, but need not, be created by a standard checkpoint utility. The auto-restart system 100 can also display the SQL statements that were being executed by the terminated child
30 processes 130, as well as any variables in use by the SQL statements. Upon request, or by

default, the auto-restart facility can additionally display actual SQL counts used by terminated processes 130.

The task control block 120 of the auto-restart system 100 can detect standard exit codes of child processes 130. For example, in an exemplary embodiment, the task control block 120 can detect the four main exit points for a COBOL program: GOBACK, STOP RUN, USER ABEND, and SYSTEM ABEND. The task control block 120 can evaluate the exit type and last SQL error code and, based on this information, can determine whether and how to rollback and restart the child process 130. Because the auto-restart system 100 can recognize standard exit codes, no modifications to individual programs need be necessary to run such programs through the auto-restart system 100.

An exemplary embodiment of the auto-restart system 100 can, but need not, utilize one or more copybooks for implementing automatic restarts. For example, and not limitation, a copybook can be used to retrieve execution parameters.

For illustrative purposes only, exemplary copybook fields are provided below. In the provided exemplary copybook fields, the abbreviation “ARF” stands for auto-restart facility, and refers to a field related to the auto-restart system 100. An exemplary embodiment of the auto-restart system 100 can include the following copybook fields:

Copybook Field	Description
:ARF:-RETURN-CODE	Internal variable used to pass return code values between the utility sub-routines.
:ARF:-VERSION	Internal variable used to mark the version of the copybook.
:ARF:-DB2-PROGRAM-NAME	Name of a program to execute as a child process, which can be called by WMDB2AR1.
:ARF:-PROGRAM-SEQ	Internal variable used to indicate which program sequence is being executed. If a multi-threading mechanism is provided, this field can allow the calling program to determine in which thread sequence it is running. This variable can be set to 1 when the multi-threaded feature is not implemented.
:ARF:-THREAD-ID	Internal variable containing the task control block address. Used to detect exit method and exit codes.
:ARF:-THREAD-ECB	Internal variable representing address of Event Control Block (“ECB”) for a called program (<i>i.e.</i> , child process). An ECB is used to wait on the

	child process. When the child process terminates, the operating system posts or triggers this ECB, and control is returned to the task control block.
:ARF:-COMPLETION-IND	Internal variable used by auto-restart system to communicate the exit type for the called program.
:ARF:-SQL-COUNT-IND	Yes/No flag used to print SQL counts upon completion of the called program.
:ARF:-CHECKPOINT-RECORDS	Number of records to be used for checkpoints.
:ARF:-CHECKPOINT-TIME	Amount of time between check points.
:ARF:-JCL-PARM-LEN	Length of the data passed to the called program. This is primarily used by the utility to pass data to the called program. The called program can either use these fields or set up a standard linkage section definition and receive data on a procedure division statement.
:ARF:-JCL-PARM-DATA	Data passed to the called program. This variable can be used by the auto-restart system to pass data to the called program. The called program can either use these fields or setup the standard linkage section definition and receive the data on the procedure division statement.
:ARF:-EXECUTION-COUNT	Represents the number of times to execute the called program in a multi-threaded application. This variable can be set to 1 when the multi-threaded feature is not implemented.
:ARF:-DUMP-COLUMN-IND	Yes/No flag indicating to dump the columns of the last SQL statement if there was an error.
:ARF:-ROLLBACK-TYPE-IND	This instructs the utility on how to perform rollbacks. This variable can be used if a child process exits with GOBACK.
:ARF:-MONITOR-SQLCODES	A list of up to fifty SQL codes to be monitored for SQL contention abends. The array also contains the number of times to retry before terminating a child process, and the length of time to wait between each retry.
:ARF:-AUTO-TRAP-IND	Yes/No flag sets up the auto trap feature of the utility. This option can reduce the code in each program that checks for contention SQL codes. With this variable, the auto-restart utility can trap contention codes and automatically rollback and restart the program.
:ARF:-ABEND-ERROR-IND	Yes/No flag used to control the exit type of the task control block.
:ARF:-IGNORE-BAD-SQL-IND	Yes/No flag to tell the auto-restart system whether to evaluate the last BAD SQL code. Processing the last BAD SQL code can detect if a contention abend occurred prior to the last SQL statement, which could mean that a rollback is desirable.

:ARF:-DSNHLLI-ADDRESSES	Internally used set of fields to communicate with the database interface.
:ARF:-GOBACK-NORMAL-EXIT	Internally used variable representing an address to store the exit code of the child process.
:ARF:-AUTO-RESTART-IND	Internally used variable to indicate whether the called program is being restarted.
:ARF:-EXCLUDE-PROGRAM-IND	Yes/No flag set when the called program is excluded from the auto-restart system, such as when found in one of the EXCLUDE-PROGRAM keywords in a parm file. When a program is <i>not</i> excluded, the auto-restart system 100 can evaluate SQL codes and attempt program restarts.

In an exemplary embodiment of the auto-restart system 100, at least a portion of the functionality of the auto-restart system 100 can be determined by the parm files. Separate parm files can be provided for global values and local values of variables. The global values can be designed to apply to the most common usages of the auto-restart system 100. When these variables are defined in the global parm file, they can represent default values for the auto-restart system 100. In contrast, when these variables are defined in the local parm file, they can override the values in the global parm file, and can apply only to programs running locally.

In an exemplary embodiment, some or all of the keywords that accept a “yes” or “no” value can also accept *y, true, t, on* or *1* for “yes” values, and *n, false, f, off* or *0* for “no” values.

Values for the following variables can be defined in the parm files:

Parm File Field	Description
CHECKPOINT-RECORDS	Number of records or limit for check point record process.
CHECKPOINT-TIME	Defines the amount of time between checkpoints.
MONITOR-SQLCODES	Identifies one or more SQL codes to consider for contention, or other resource unavailability, to evaluate for auto-restarts. This keyword can be entered multiple times to define each of the SQL codes to monitor.
SHOW-SQL-COUNTS	Yes/No value determines whether to display the SQL counts at the end of successful completion of a called program.
DUMP-COLUMNS	Yes/No value to request the auto-restart system to dump the SQL statement, any associated host variables, and any values of those variables if a called program fails with an unexpected SQL code.
ROLLBACK-FORMAT	This specifies how the ARF utility will perform rollbacks.

<p>AUTO-TRAP-SQLCODES</p>	<p>Yes/No flag sets the auto-trap functionality. When this is activated, the monitored SQL codes can be evaluated after every SQL statement.</p> <p>When a monitored code is detected, the called program can be terminated with a user abend, and control is not returned to the called program. The user abend can force an automatic rollback of the current transaction. The auto-restart system can then restart the called program.</p>
<p>ABEND-ON-ERROR</p>	<p>Yes/No flag determines how the auto-restart system will exit on non-zero return codes. The global setting can be NO. A value of YES can abend the current step and process the error disposition option for all files.</p>
<p>IGNORE-BAD-SQL</p>	<p>Yes/No flag determines whether auto-restart system evaluates the last bad SQL code. The last bad SQL code could be the last SQL statement executed or could be from a past call.</p> <p>When this is set to NO, the auto-restart system can evaluate only the last bad SQL code when a rollback is issued by the called program.</p> <p>This option is useful with programs that handle their own rollbacks. If some programs auto-correct their own errors and continue processing, it may not be desirable for the auto-restart system to attempt to rollback and restart these programs.</p>
<p>EXCLUDE-PROGRAM</p>	<p>This keyword expects a program name for a program to be excluded from the auto-restart system.</p>

Conclusion

The above implementation is provided for illustrative purposes only, and does not limit the scope of embodiments of the auto-restart systems 100. Various other implementations could be used to employ an exemplary embodiment of the system 100.

Accordingly, as described above, exemplary embodiments of the auto-restart system 100 and method can be used to monitor database processes, and to restart such processes after they terminate abnormally.

While embodiments of the auto-restart systems and methods have been disclosed in some exemplary forms, it will be apparent to those skilled in the art that many modifications,

additions, and deletions can be made without departing from the spirit and scope of the auto-restart systems, methods, and their equivalents, as set forth in the following claims.

CLAIMS

1. A process management system comprising:
 - a first task control block running in a database management system;
 - a computer processor for executing operations of the first task control block;
 - a first child process executing a first set of instructions in a database management system, the first child process outputting an exit code and terminating upon encountering an unresolved database contention error while executing the program code; and
 - the first task control block being configured to detect the exit code and at least a portion of the first set of instructions, and automatically restart the first child process according to one or more parameters based on the exit code and the detected portion of the first set of instructions.
2. The process management system of Claim 1, the first task control block restarting the first child process without human intervention.
3. The process management system of Claim 1, the first task control block determining whether to rollback the first child process based on the exit code output by the first child process.
4. The process management system of Claim 1, the first task control block determining when to restart the first child process based on at least one of the exit code and the detected portion of the first set of instructions.
5. The process management system of Claim 1, the first task control block rolling back the first child process to a checkpoint before restarting the first child process.
6. The process management system of Claim 1, the first task control block being in a hung state until for the first child process terminates.
7. The process management system of Claim 1, further comprising a second child process for running a second set of instructions in the database management system.

8. The process management system of Claim 7, the second child process being spawned by a second task control block.
9. The process management system of Claim 7, the second child process locking a shared resource required for the first child process to execute the first set of instructions, thereby causing the database contention error.
10. A process management system comprising:
 - a plurality of parent processes, including a first parent process and a second parent process;
 - a computer processor for initializing the plurality of parent processes;
 - a plurality of child processes, including a first child process spawned by the first parent process, and a second child process spawned by the second parent process;
 - a shared resource accessible by each of the plurality of child processes, the shared resource being locked by the first child process;
 - the second child process encountering an error upon attempting to access the locked shared resource and, upon encountering the error, outputting an exit code and terminating; and
 - the second parent process detecting the termination of the second child process, and automatically restarting the second child process if the second child process meets one or more restart conditions.
11. The process management system of Claim 10, the shared resource being a database record.
12. The process management system of Claim 10, the second parent process determining how to restart the second child process based at least partially on the exit code.
13. The process management system of Claim 10, the second parent process determining when to restart the second child process based at least partially on the exit code.

14. The process management system of Claim 10, the second parent process determining whether to restart the second child process based on a command that the second child process attempts to execute before its termination.
15. The process management system of Claim 10, further comprising a participation list, wherein the second parent process compares one or more characteristics of the second child process to the participation list to determine whether to restart the second child process.
16. The process management system of Claim 10, the second child process waiting for the shared resource for a predetermined time period before being terminated.
17. The process management system of Claim 10, the second child process checking availability of the shared resource a predetermined number of times before being terminated.
18. The process management system of Claim 10, the second child being rolled back to a checkpoint.
19. A computer program product embodied in a computer-readable medium, the computer program product comprising an algorithm adapted to effectuate a method for managing processes, the method comprising:
 - providing a plurality of task control blocks;
 - configuring each of the plurality of task control blocks to initialize a corresponding child process for executing program code;
 - detecting an abnormal termination of a first child process corresponding to a first task control block; and
 - determining whether to restart the first child process automatically based on the program code executed by the first child process.
20. The computer program product of Claim 19, the method further comprising configuring the first task control block to restart the first child process automatically.

21. The computer program product of Claim 20, the method further comprising configuring the first task control block to roll back the first child process to a checkpoint before restarting the first child process.
22. The computer program product of Claim 19, the method further comprising configuring the first task control block to compare the program code against a participation list to determine whether to restart the first child process.
23. The computer program product of Claim 19, the first child process executing the program code in a database management system.
24. The computer program product of Claim 19, the abnormal termination of the first child process caused by a database contention error.

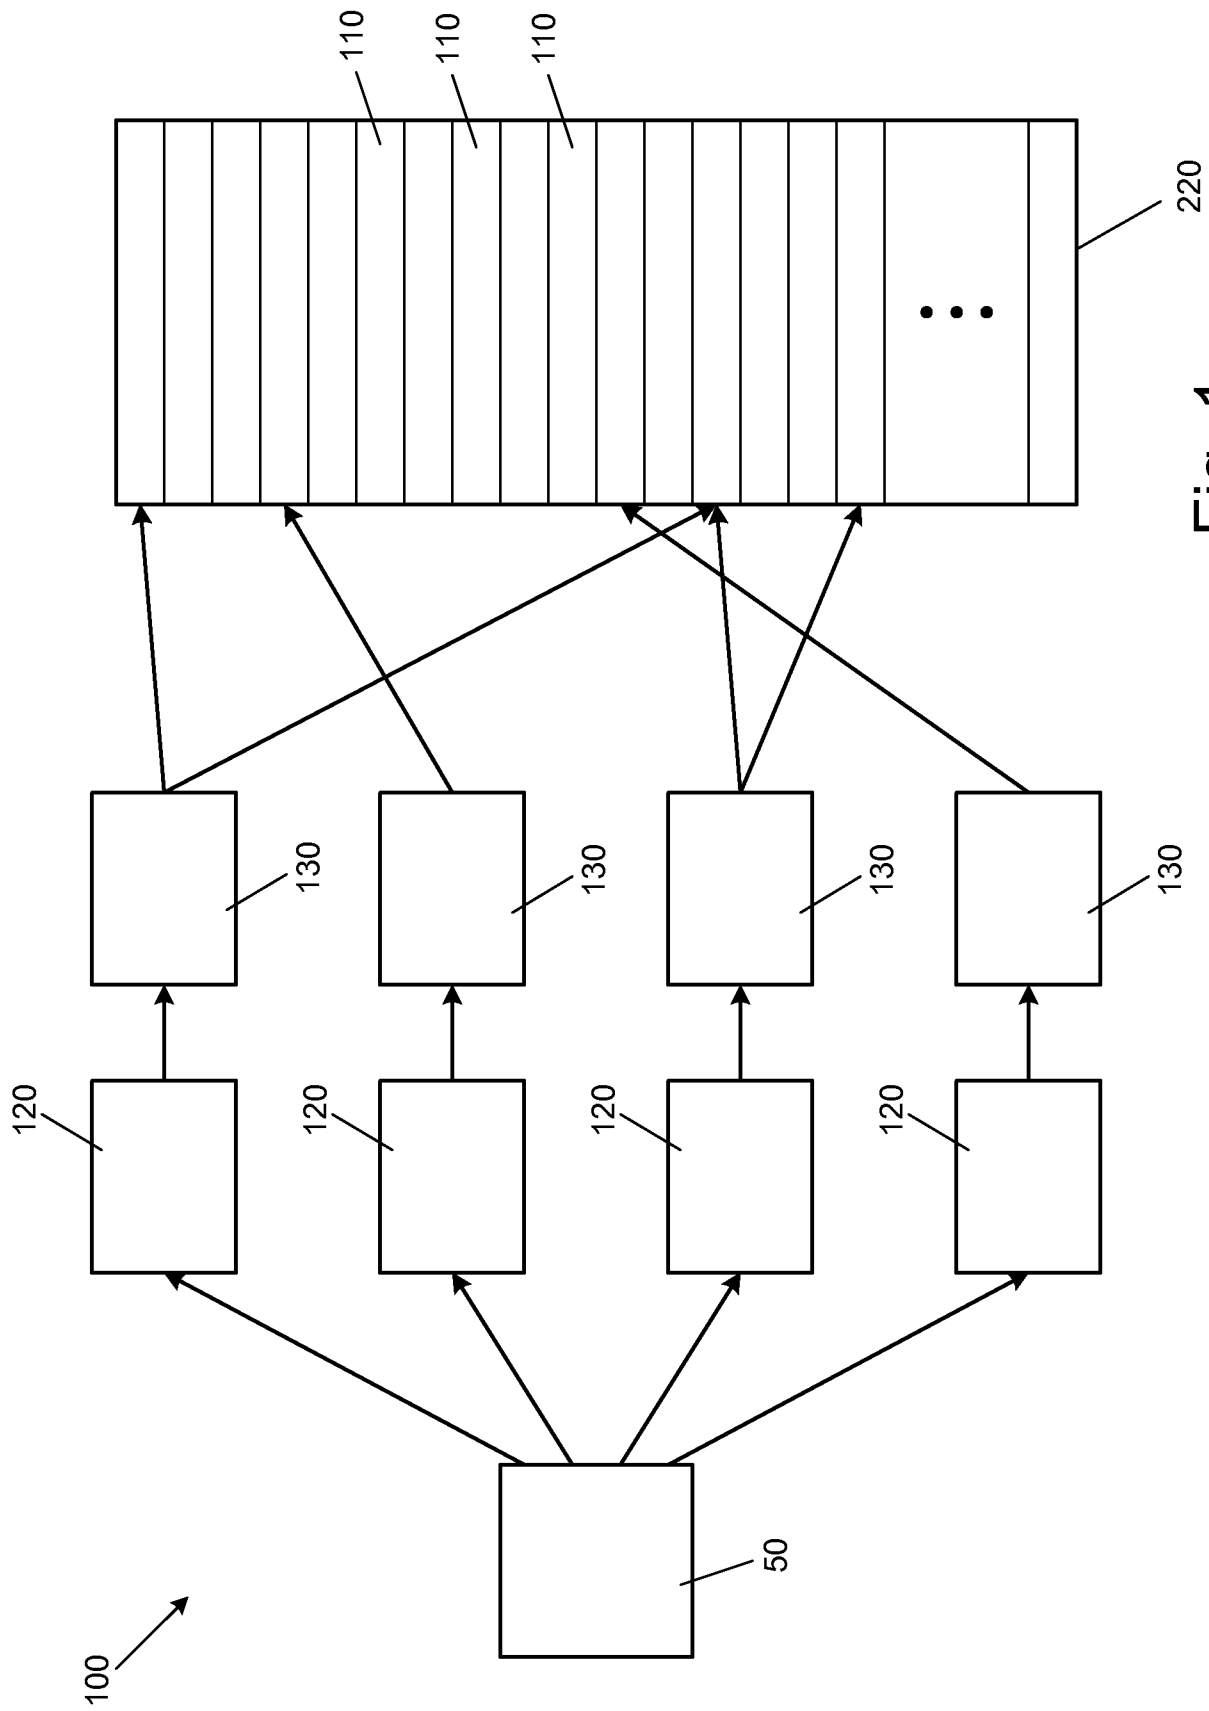


Fig. 1




Fig. 2
Prior Art




Fig. 3

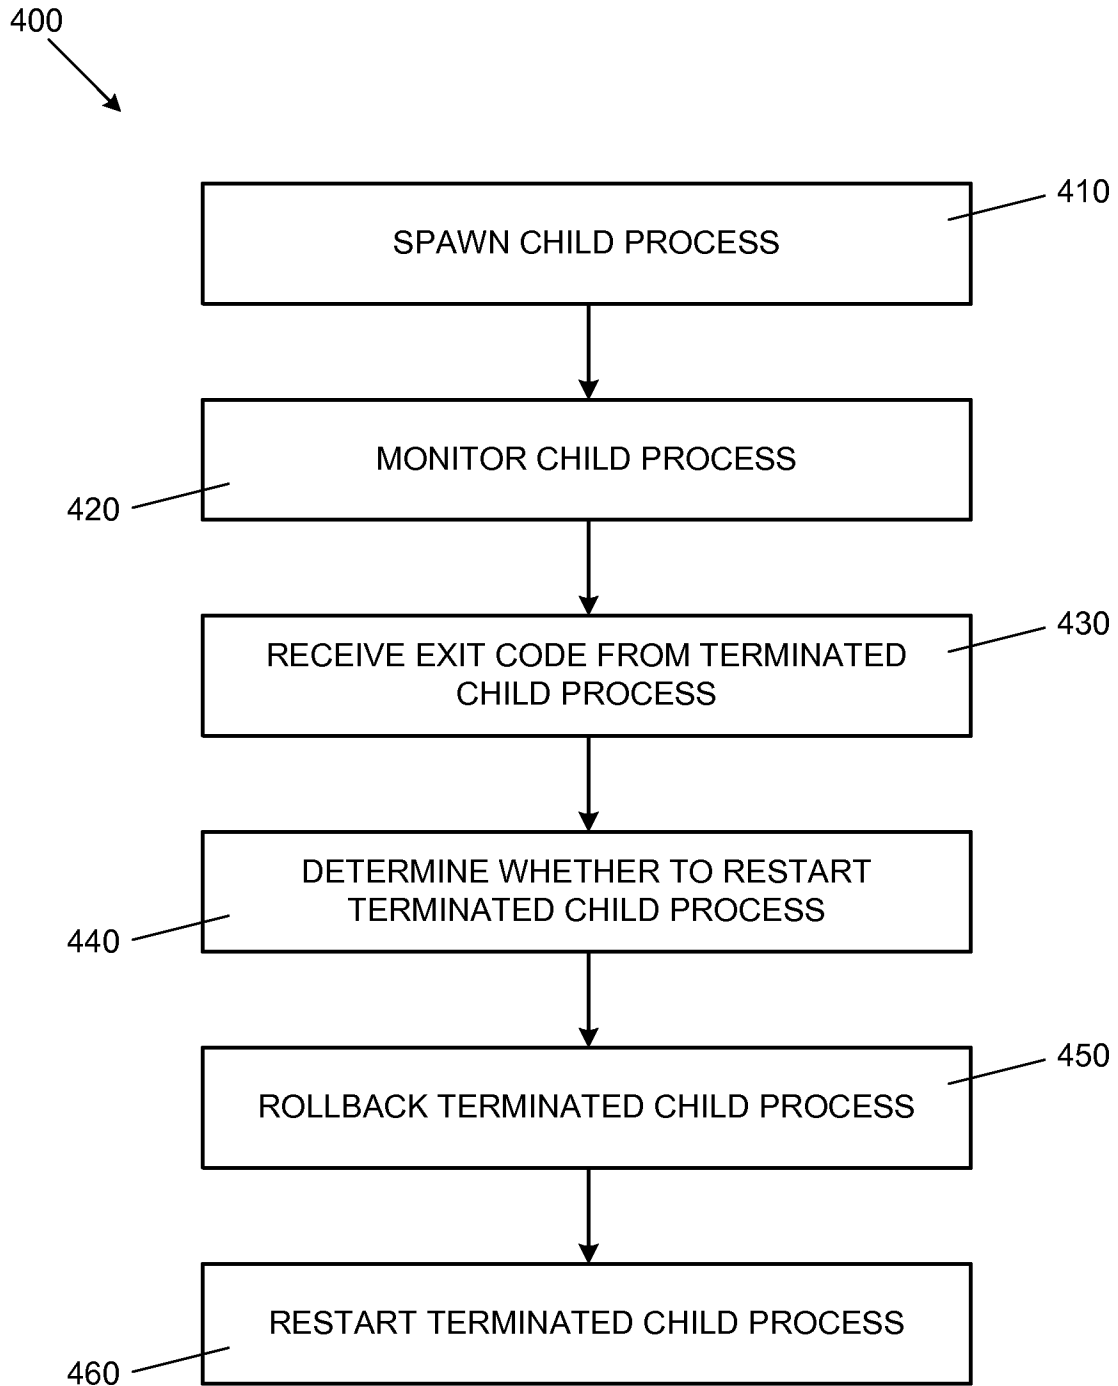


Fig. 4

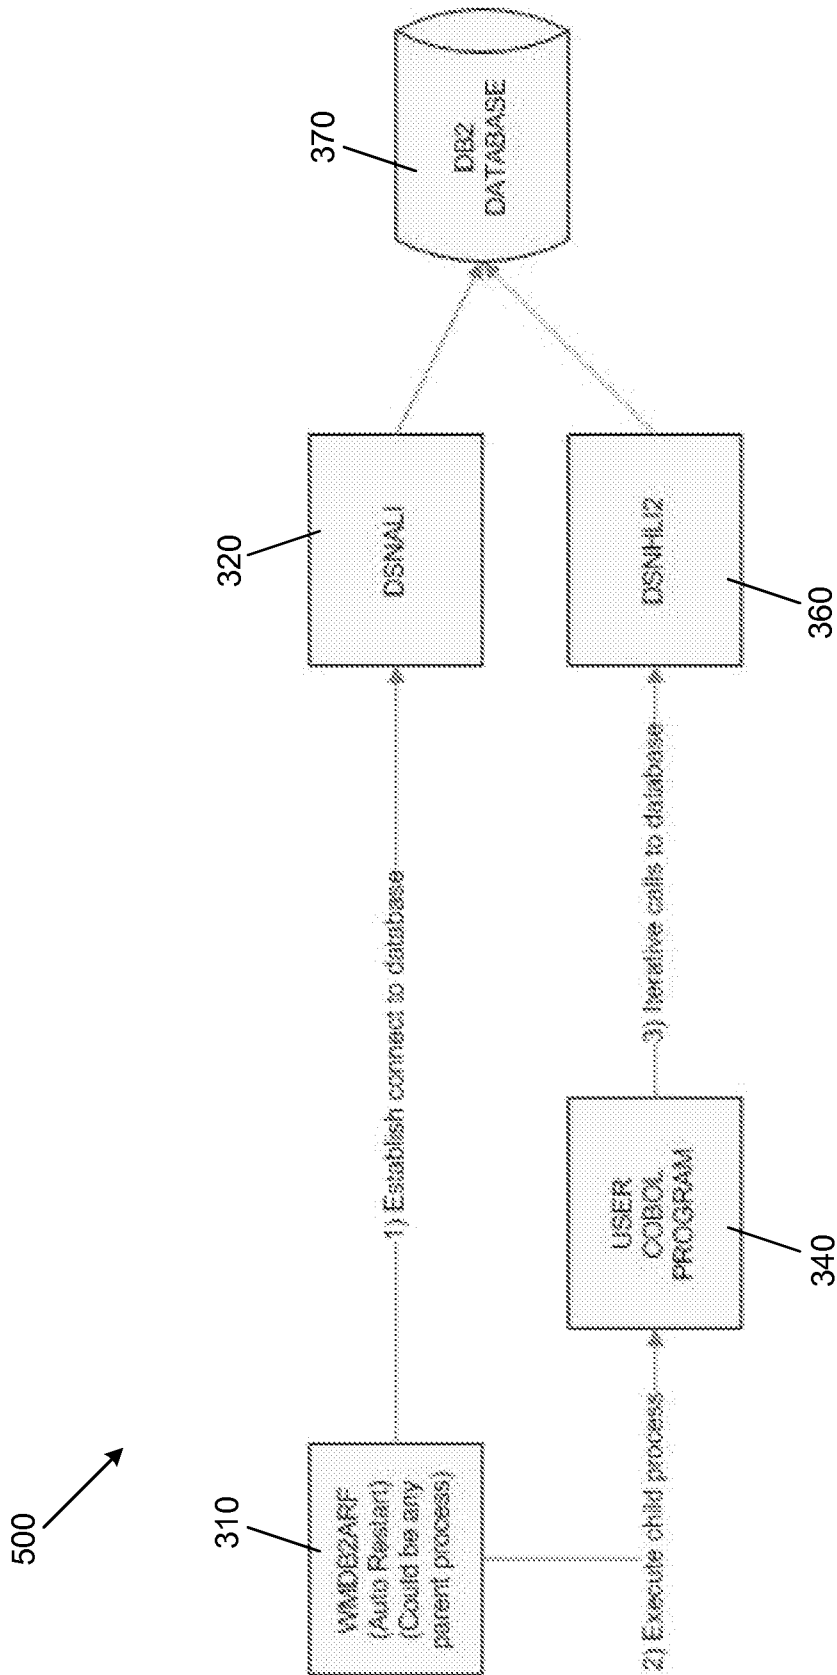


Fig. 5




Fig. 6
Prior Art